KGI IWM

The Strategic Playbook Series

Global AI Asset Allocation: Unlocking the Investment Framework for Future Growth

29 June 2026

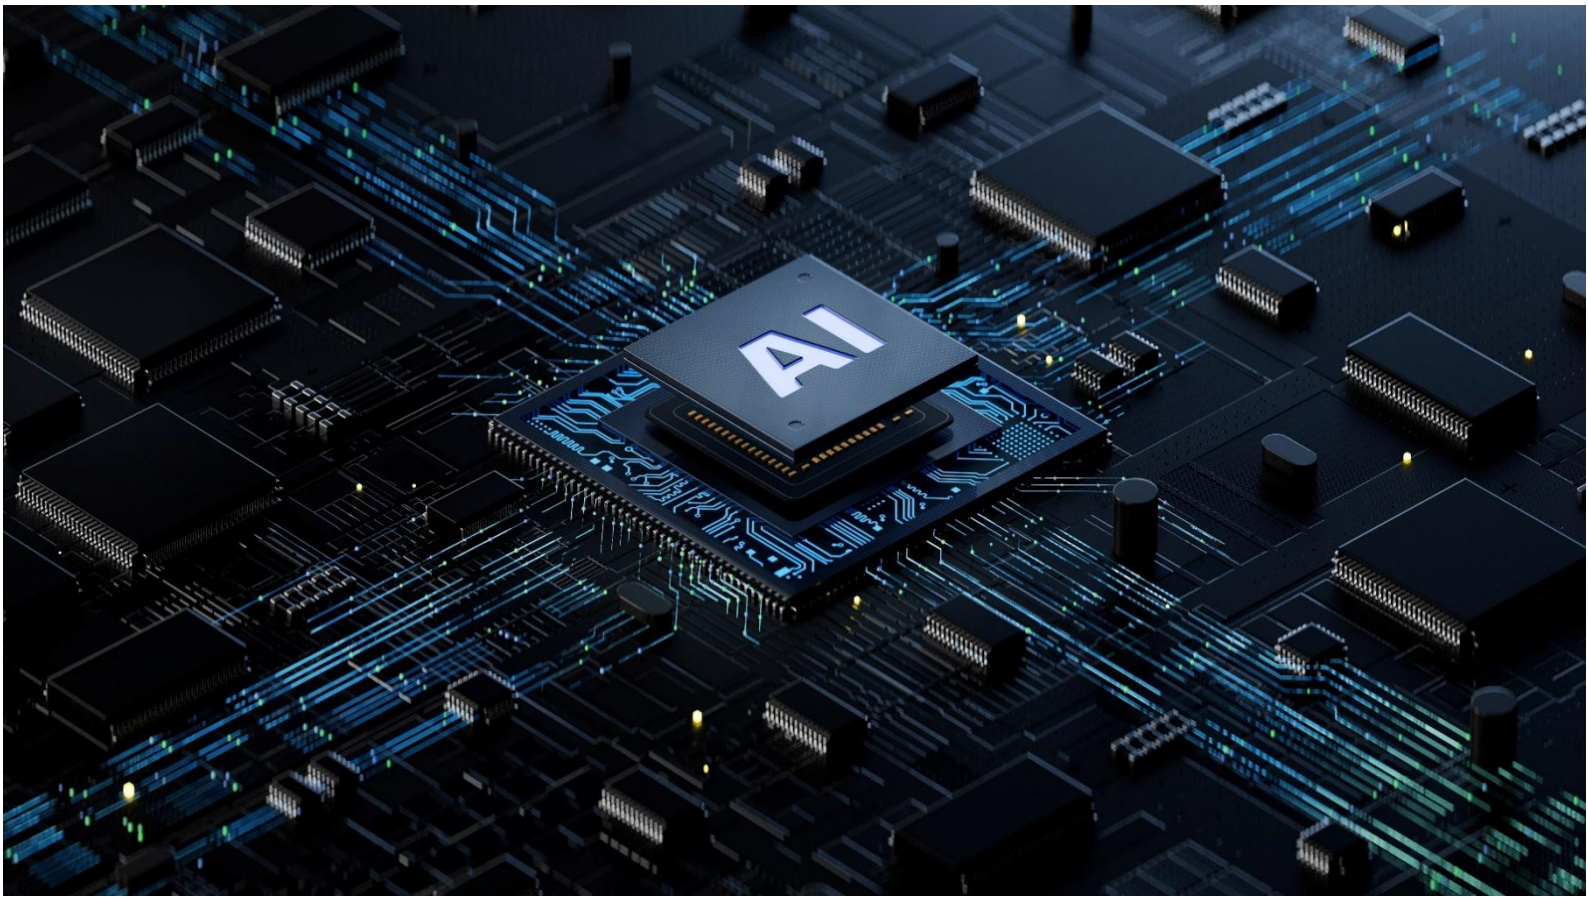

Table of Contents

Investment Products and Solutions

Kevin Yin
Equity Advisory
kevin.h.yin@kgi.com
+852 2878 6798
SFC CE: AJM689

Raymond Mok, CFA
Equity Advisory
raymond.mok@kgi.com
+852 2878 4663
SFC CE: BHJ465

Derek Yip, CFA
Equity Advisory
derek.yip@kgi.com
+852 2878 4263
SFC CE: BSQ196

Cusson Leung, CFA
CIO
cusson.kt.leung@kgi.com
+852 2878 6998
SFC CE: ADU276

EXECUTIVE SUMMARY

This report systematically analyzes **the seven critical issues** currently dominating the investment landscape, built upon a deep deconstruction of **the 20 core layers** within the global AI ecosystem. By examining these topics through the pivotal lenses of infrastructure, technological monopolies, monetization, valuation boundaries, and investment risks, we aim to construct **a forward-looking investment framework designed for both growth and resilience.** Backed by this rigorous fundamental research, we are pleased to introduce **the KGI Global AI-30 Basket** investment solution.

1. The Four Strategic Pillars

In the global AI ecosystem, Power sets the hosting capacity, Foundry determines compute supply, Network/Storage drives compute efficiency, and LLM/Agent ultimately determines monetization profits. In asset allocation, leaders across these four sub-sectors exhibit strong pricing power and defensiveness.

Global AI Ecosystem: A Panoramic View of the Four Strategic Pillars

Strategic Segments	Strategic Importance	Underlying logic	Core Companies
1. Energy & Nuclear / Power Infrastructure	**No power?** The absolute survival baseline for compute expansion; power supply directly determines whether a data center can be built.	Hyperscale data centers consume as much power as a mid-sized city. Zero-carbon, 24/7 firm nuclear power and Small Modular Reactors (SMR) have become scarce resources. Tech giants are signing multi-decade contracts to lock in grid capacity ahead of others.	Constellation Energy (CEG), GE Vernova (GEV), NextEra Energy (NEE)
2. Wafer Foundry and Advanced Packaging	**No chip?** The widest moat in the global tech supply chain; Can chips be made?	Nearly 100% of top-tier AI accelerator chips—whether from NVIDIA, AMD, or tech giants' ASICs—exclusively rely on TSMC (TSM US / 2330 TT) for advanced nodes and CoWoS packaging. If geopolitics or capacity bottlenecks erupt, the global AI compute build-out would hit an immediate physical standstill.	TSMC (TSM US / 2330 TT)
3. High Bandwidth Memory and High-Speed Networking	**Bandwidth**—determines whether chips can run at high speed. The biggest bottleneck today.	Without HBM (high-bandwidth memory), a GPU is just a data-starved shell; without high-speed interconnect chips and switch architectures, tens of thousands of chips cannot be integrated into a supercomputer.	SK Hynix (000660 KS), Micron (MU), Broadcom (AVGO)
4. Core foundation models and Agent ecosystem	**Applications?** The ultimate traffic gateway and ecosystem moat;	Frontier foundation models are rapidly evolving into multi-agent systems (MAS), directly taking over enterprise productivity workflows (e.g., billed by Token throughput). This will set the monetization ceiling for the entire AI industry. Once enterprise (B2B) clients are deeply integrated, switching costs become prohibitive, effectively handing over the master switch for their workflows and financial data.	Anthropic, OpenAI, Microsoft (MSFT) (Copilot)

Source: KGI

2. The Top-Five Global Leaders

ASML (ASML) (exclusive supplier of EUV lithography tools)—> TSMC (TSM US / 2330 TT) (exclusive foundry and advanced packaging)—> NVIDIA (NVDA) (compute leader) / Broadcom (AVGO) (networking leader) / AMD (AMD) (NVIDIA rival)

The five undisputed giants of the global AI ecosystem — highly interdependent with intense co-opetition

Big 5	Industry Roles	Core Moat (Moat)	Core business	Profitability
NVIDIA (NVDA)	Fabless (system vendors); the undisputed leader in today's AI compute market	**Irreplaceable GPU** hardware assets and **CUDA** software ecosystem.	Full-Stack AI Computing Platform (chips, NVLink network, server systems, CUDA software)	Windfall-profit king, GM up to 75%.
TSMC (TSM US / 2330 TT) (TSMC)	Foundry (wafer fab); the most defensive core asset in the compute supply chain	**Exclusive control** of over 99% of global top-tier **AI chip foundry and advanced packaging** (NVIDIA (NVDA), AMD (AMD), Apple (AAPL), etc.).	Leading-edge logic chip manufacturing (3nm/2nm) and advanced packaging (CoWoS, SoIC)	Earnings profile is highly resilient, with GM at 66.2%.
ASML	Hardware (semiconductor equipment); benefiting from capacity bottlenecks, the order book is already fully booked.	**Sole supplier** of the ultimate strategic equipment essential for manufacturing advanced AI chips (**EUV lithography machines**).	Lithography systems (EUV, High-NA EUV)	GM ~50%, NM ~30%.
Broadcom (AVGO)	Fabless (customization/networking); core beneficiary of network architecture upgrades	Core Ethernet switch chips that **absolutely dominate** 100,000-card-class hyperscale data centers.	Custom AI ASICs (e.g., Alphabet (GOOGL) TPU, Meta (META) accelerators) and high-bandwidth network switch chips (Tomahawk series)	Exceptionally strong profitability; GM at 65%–70%.
AMD	Fabless (system vendors); agile potential challengers	The industry's only high price-performance alternative capable of **competing head-on** with NVIDIA (NVDA)'s flagship GPU.	Compute alternative solutions (MI300/ MI400 series accelerators) and x86 server CPU	Margins are currently below peers but are expected to expand.

Source: KGI

3. The Four Mega-bottlenecks

Four core Mega-Bottlenecks across the 20 layers of the global AI ecosystem

Bottleneck level	Essence of the Bottleneck	Approaches to Overcome	Time to Master	Biggest risk
Level 1 / Level 9 (Energy & Power / Transmission and Distribution Equipment)	**We have nuclear power plants, but electricity can't be transmitted to the compute center (physical equipment shortage).**	Bypass the grid with self-built microgrids; bring in overseas supply chains.	2029 - 2031	Supply chain disruptions caused by trade protectionism
Layer 5 / Layer 7 (Packaging / Storage)	**Wafers are sufficient, but packaging throughput and memory production constrain compute shipments (CoWoS and HBM yield bottlenecks)**	TSMC (TSM US / 2330 TT) packaging capacity to expand several-fold; ASML (ASML) equipment skewing toward memory	2027	Geopolitical risk in the Taiwan Strait
Layer 17 (Data Platform)	**High-quality web text is running out; internal enterprise data is disorganized.**	High-fidelity synthetic data, vector knowledge graph plugin	2028 - 2029	Model inbreeding-induced degradation, data copyright litigation
Layer 20 (Governance and Regulation)	**EU AI Act hard-lands in August 2026; black-box models hit a compliance brick wall**	Commercializing mechanistic interpretability tools and compliance gateways	2028	Compliance shock leads to widespread delays in enterprise AI applications

Source: KGI

4. The Breakdown of Global USD 7.6 trillion Capex

The global AI industry is in the largest infrastructure build-out supercycle in business history. From 2026 to 2031, cumulative AI-related Capex is expected to reach USD 7.6 trillion. In 2026 alone, the combined Capex of the top five cloud giants exceeds USD 600 billion. At its core, today's hundreds of billions in Capex reflect tech giants deploying high-margin cash from current software businesses to lock in the next 10 years of global physical infrastructure (power grids, nuclear power, data center land).

Capex by business segment **Capex by country**

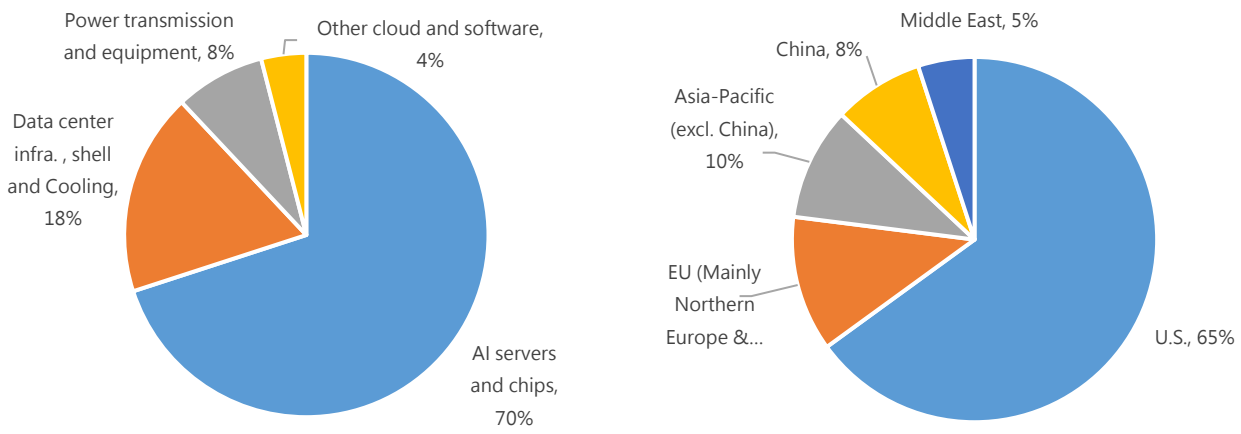

Source: KGI

5. The Profitability Trend

In 2026, the global AI ecosystem's profit distribution exhibits an extremely polarized barbell structure.

Windfall-Profit Chokepoints	Key representatives	Earnings Performance (2026)	Underlying logic
1 Proprietary AI silicon and IP layer	NVIDIA (NVDA) ARM	GM: ~74.9%; OM: ~65.6%; NM: ~71.4%	Pricing power: companies are buying chips, but more importantly the compute enabled by the CUDA software ecosystem. As R&D is amortized over scale, profits exhibit nonlinear expansion.
2 Monopolistic AI Gatekeepers: Storage Layer	SK Hynix (000660 KS), Micron (MU) Samsung Electronics (005930 KS)	GM: 53%-58%; OM: 35%- 42%; NM: 28%-34%	Must-have demand: HBM4 is shifting from a traditional commodity to a highly customized, ASIC process. Major buyers are prepaying 18 months in advance to lock in capacity, ending price wars and lifting margins.
3 High-Moat B2B SaaS and AI Assistants	Microsoft (MSFT), Palantir (PLTR)	GM: 70%- 80%; OM: 40%- 45%; NM: 32%-36%	Extremely low marginal cost: Microsoft (MSFT) is deeply embedding AI into the Office 365 ecosystem, with the underlying compute on Azure under its own control. The marginal cost of upselling enterprise subscriptions is minimal, yielding high net profit.
Loss drivers	**Key players**	**Earnings Performance (2026)**	**Underlying logic**
1 Pure AI standalone LLM layer	OpenAI, Anthropic, AI startups	GM: 25%-40%; OM: deeply negative (<-50%); NM: severe losses	The "sandwiched layer" squeezed from both ends: upstream, forced to pay cloud giants hefty compute rent (in effect paying for NVIDIA (NVDA)); downstream, hit by a global API cut-throat price war, with revenue unable to cover compute and talent costs.
2 Asset-heavy, highly levered second-tier compute clouds	CoreWeave; Lambda Labs	GM: 65.5%; OM: -6.9%; NM: -35.6%	Eroded by depreciation and interest: relies on debt financing collateralized by GPUs to acquire assets. While the reported GM on compute rentals appears high, once online it faces hefty equipment depreciation, interest, and rent, pushing net income deep into the red.
3 Low-Barrier AI Server Assembly & Integration	Super Micro Computer (SMCI) Foxconn Industrial Internet (601138 CH)	GM: 6.3%- 9.5%; OM: 3%- 5%; NM: 1.5%- 3%	Heavy-lifting without core IP: lacks proprietary chip IP, only doing hardware contract manufacturing and assembly. Amid cutthroat competition, forced to slash margins to win orders—resulting in revenue over 100 billion with paper-thin net profit.

Note: GM = gross margin; OM = operating margin; NM = net margin. Financial data reflects the latest 2026 developments and actual disclosures.

Source: KGI

Global AI Industry Profit Margin Trends, 2026–2030

Core trends	Evolution Roadmap and Timeline	Key Catalysts (Catalysts)	2028-2030 Earnings Outlook
1 Upstream Chip/IP Layer: Windfall Profits Have Peaked and Are Rolling Over	2026 to hit the windfall-profit peak; expected to revert to rationally elevated levels after 2028.	1. LLMs are largely through the infrastructure-build training phase; compute focus is shifting to inference. 2. Cloud hyperscalers' in-house ASICs (e.g., Alphabet (GOOGL) TPU, Amazon (AMZN) Trainium 2) and AMD (AMD) are ramping at scale, breaking the monopoly.	Upstream hardware GMs for players such as NVIDIA (NVDA) are expected to correct from the 75% physical ceiling toward a 60%–65% rational range, with premiums ceded to downstream players.
2 Compute leasing market sees a brutal shakeout	Bankruptcy wave in 2026-2027; after the industry shakeout, surviving leaders will swing to profitability.	Highly leveraged, debt-funded second-tier compute cloud providers will face a cash flow crunch as equipment depreciation costs hit all at once, ushering in a "darkness before dawn" and a capacity shakeout in the market.	For the surviving leading compute-cloud providers, after equipment depreciation rolls off (post-2028), GM will rebound from the 2026 trough of 8%-15% to 15%-25%, and NM will swing to positive.
3 Ultimate profit harvesters will fully shift to the application layer	2028–2030 will usher in a golden harvest phase, with the profit pool completing its shift from pick-and-shovel sellers to power users.	Underlying compute costs have collapsed by 90% as technology iterates, materially reducing software vendors' operating costs. Vendors that control end-to-end enterprise workflows and enjoy structural B2B stickiness command outsized premium pricing power.	Vertical AI and SaaS vendors will enjoy ultra-low compute costs and premium subscription pricing, with net margins poised for a nonlinear breakout to 35%+.

Source: KGI

6. AI Business Model Comparison: US vs China

- **US:** Leveraging the world's deepest USD capital pool, a best-in-class semiconductor supply chain, and brute-force supercomputing power, it continues to go big in pushing the frontiers of AGI (general artificial intelligence) and multimodal ecosystems.

- **China:** Under chip constraints and without an absolute capital advantage, it has driven AI inference costs to near-free by leveraging a formidable ability to wring performance from algorithms, relentless engineering optimization, and a highly pragmatic free/open-source strategy, showing strong aggressiveness in productivity democratization and industrial deployment.

7. Has the Global AI Trade Peaked Out? Not Yet

1) **Today's valuation of 37x vs 45–100x during bubble periods**. In past bubbles, P/E in the affected segments at least surged to 45–72x on a trailing 12-month basis. In Asia in particular, Taiwan and Japan reached peak trailing 12-month P/E of 72x and 100x in 1989. By contrast, the MSCI ACWI Semiconductors sector now trades at 37x trailing P/E.

2) **Fundamentals remain robust with earnings estimates keep moving higher.** Over the past 3 months, corporate EPS revisions have been raised 10% more than the broad market. 12-month forward EPS growth is estimated at 40%, versus 12% for the MSCI AC World Index.

3) **Upstream capacity constraints will take years to resolve.** Supply remains tight because: (A) Long fab build cycles (it takes 3 years to build a wafer fab); (B) TSMC (TSM US / 2330 TT) maintains disciplined Capex management (Capex is only up ~80% vs the pre-GenAI period); (C) Memory capacity can only increase ~30% across 2026–2027; and (D) Capex growth at TSMC and memory makers lags hyperscalers materially

Relative to the TMT bubble, overinvestment is limited: IP equipment and software spending as a share of GDP has risen 17% since 2022, versus a 42% increase during 1994–2000.

4) **TAM expectations have significant room for upward revision.** Assuming semiconductor hardware and software account for 2% of global GDP in 2030, we estimate industry revenue (TAM) at USD 3 trillion. If the 2030 net margin stays at the 2025 level of 35%, the implied P/E is 16x; if the 2030 net margin is steady at 25% (vs. 25% in 2025), the implied P/E is 22x. From this perspective, neither TAM assumptions nor valuations look overly aggressive. As a reference, oil's share of global GDP has averaged 3% historically (spent long stretches between 4% and 5%, peaking at 10%), and AI may warrant a higher share.

5) **Sheer market cap has already capped how far investors can Overweight.** Many funds stipulate that a single-stock position cannot exceed 10% of total portfolio market value. This means these investors now struggle even to reach a Market-weight allocation versus the index. TSMC (TSM US / 2330 TT), Samsung Electronics (005930 KS) and NVIDIA (NVDA) are already facing this issue. The problem is even more acute in Growth mandates.

6) **Future deployable capital equals 60% of the current sector market cap.** About USD 2.3 trillion is flowing into the US equity market via dividends, share buybacks, and cash-funded M&A. Together with up to USD 7.8 trillion still parked in money market funds, the total pool exceeds USD 10 trillion—equivalent to 60% of the sector's current market cap of USD 16.6 trillion.

Market cap size limits how much investors can Overweight.

Stocks	Index	Market Cap (%)
Broadcom	MSCI US	3.5%
ASML	MSCI Europe	4.9%
NVIDIA	MSCI US	8.0%
SK Hynix	MSCI EM	8.7%
ASML	MSCI Europe growth	9.6%
NVIDIA	MSCI ACWI growth	10.0%
SEC	MSCI EM	10.8%
TSMC	MSCI EM	15.6%

Source: UBS, Refinitiv DataStream, KGI

8. How to Position? KGI Global AI-30 Basket

KGI Global AI-30 Basket - Investment Thesis

Listed companies breakdown: US: 19; China: 5; Taiwan: 2; South Korea: 2; Netherlands: 1; Japan: 1

Name and Ticker	Investment Thesis	Market Cap (USD mil)	Sales YoY (%) 2026	Sales YoY (%) 2027	EPS YoY (%) 2026	EPS YoY (%) 2027	EBIT Margin 2026	EBIT Margin 2027	ROE 2026	ROE 2027	FCF Yield	EV/Sales	P/E 2026	P/E 2027
(1) Global AI Infra Monopoly														
Nvidia (NVDA.US)	Monopolistic GPU/CUDA; compute dominance; high margins, reasonable valuation.	4,736,908	65.5	81.5	**59.5**	**88.7**	63.4	66.2	92.9	92.6	**2.5**	18.3	41.8	**21.8**
TSMC (TSM.US)	Monopolizes top-tier global wafer fabrication and packaging; high margins, reasonable valuation.	2,256,064	36.5	27.8	**49.2**	**26.7**	57.3	56.9	38.8	37.1	**1.9**	14.3	28.1	**22.1**
Broadcom (AVGO.US)	Monopolizes high-speed networking chips; asset-light, high gross margin, strong free cash flow.	1,802,695	65.6	63.8	**69.4**	**66.2**	65.5	65.0	54.6	59.6	**1.8**	24.5	32.8	**19.7**
Micron (MU.US)	HBM3e/HBM4 breakthroughs; top share contender; geopolitical tailwinds.	1,370,586	238.2	89.8	**764.9**	**107.1**	74.6	81.6	74.5	71.8	**1.9**	14.9	16.9	**8.2**
Samsung Electronics (005930.KS)	Downstream vertical integration across the full value chain; early HBM misjudgment; catching up.	1,251,908	106.6	27.0	**495.6**	**38.5**	51.5	57.1	49.8	44.1	**2.9**	4.7	7.5	**5.4**
SK Hynix (000660.KS)	Global HBM share No.1 (58~61%); holds the vast majority of NVIDIA (NVDA) core orders.	1,241,648	251.9	47.4	**394.9**	**44.8**	76.5	77.8	92.8	62.7	**2.2**	14.3	8.8	**6.1**
ASML (ASML.NA)	Monopolizes EUV lithography tools; relatively oversold; earnings revisions	702,852	19.6	23.2	**28.2**	**33.5**	36.6	40.1	55.5	58.9	**1.4**	18.2	50.3	**37.7**
	Simple Average		**112.0**	**51.5**	**265.9**	**57.9**	**60.8**	**63.5**	**65.6**	**61.0**	**2.1**	**15.6**	**26.6**	**17.3**
	Market Cap Weighted Average		**97.1**	**59.5**	**201.8**	**65.4**	**62.5**	**64.9**	**70.7**	**67.6**	**2.2**	**16.4**	**29.9**	**18.0**
(2) Mega 7														
Nvidia (NVDA.US)	Monopolistic GPU/CUDA; compute dominance; high margins, reasonable valuation.	4,736,908	65.5	81.5	**59.5**	**88.7**	63.4	66.2	92.9	92.6	**2.5**	18.3	41.8	**21.8**
Alphabet (GOOGL.US)	GCP cloud; Gemini best-in-class; search engine and ads monetized	4,175,544	23.2	21.2	**11.1**	**7.4**	40.3	40.6	49.5	38.0	**1.5**	9.8	24.1	**22.4**
Apple (AAPL.US)	AI applications at scale; near-term hardware innovation slowing; resilient free cash flow.	4,041,226	14.9	9.0	**17.5**	**10.2**	32.4	31.9	132.5	108.0	**3.2**	8.8	31.4	**28.5**
Microsoft (MSFT.US)	Azure cloud, OpenAI; generative AI embedded into enterprise software; compounding growth.	2,620,975	16.9	16.7	**25.5**	**13.9**	46.6	46.7	31.1	27.6	**2.8**	8.4	20.6	**18.1**
Amazon (AMZN.US)	AWS cloud growing strongly; e-commerce margin expansion; balanced offense and defense.	2,441,971	15.1	13.2	**14.9**	**12.5**	12.6	14.1	19.8	17.8	**-0.1**	3.4	22.0	**19.6**
Tesla (TSLA.US)	Full self-driving FSD, energy storage and humanoid robots; highest beta.	1,408,847	8.2	15.2	**12.8**	**30.0**	5.3	6.7	6.4	7.6	**0.5**	14.1	200.3	**154.1**
Meta Platforms	AI boosts FB/IG ad targeting and monetization; new-tech acquisition	1,378,034	25.9	19.2	**27.9**	**-4.1**	35.7	34.9	32.1	25.5	**3.5**	6.4	13.5	**14.1**
	Simple Average		**24.2**	**25.2**	**24.2**	**22.7**	**33.7**	**34.4**	**52.0**	**45.3**	**2.0**	**9.9**	**50.5**	**39.8**
	Market Cap Weighted Average		**28.6**	**30.5**	**26.9**	**28.6**	**38.9**	**39.7**	**65.6**	**57.5**	**2.1**	**10.7**	**40.1**	**30.9**
(3) Segment Leaders														
AMD (AMD.US)	Direct competitor to NVIDIA (NVDA); MI300/400 series offers superior price/performance.	868,409	43.7	53.8	**73.8**	**81.0**	26.8	31.7	14.2	21.5	**1.0**	23.0	73.5	**40.6**
Renesas (6723 JP)	Automotive chip giant; AI edge computing; smart-vehicle cycle recovery.	59,093	16.9	9.8	**37.3**	**5.6**	24.7	23.0	14.6	12.2	**3.7**	7.5	20.4	**19.3**
Sandisk (SNDK.US)	Pure-play NAND flash and enterprise SSD exposure; storage supercycle.	345,790	169.0	130.5	**2,095.5**	**191.3**	59.9	78.4	60.5	71.0	**1.3**	26.0	35.6	**12.2**
Marvell (MRVL.US)	Absolute leader in Optical/DSP; core partner on ASICs; high-margin, asset-	246,046	42.1	40.6	**80.9**	**42.4**	35.3	36.3	2.8	12.0	**0.7**	28.4	99.1	**69.6**
Coherent (COHR.US)	Global optical-communication materials giant; asset-heavy with high depreciation; profits surge once mass production ramps.	79,674	21.5	36.2	**54.5**	**52.6**	20.4	22.4	9.9	12.9	**-0.7**	12.3	74.7	**48.9**
Lumentum (LITE.US)	Monopolizes global EMI optical chips; order visibility directly tied to total global compute expansion.	67,061	82.2	85.7	**295.9**	**121.6**	29.1	39.0	32.7	44.8	**0.2**	27.0	105.7	**47.7**
GE Vernova (GEV.US)	Monopolizes gas turbines and grid equipment markets; high-certainty energy exposure with strong orders.	291,687	19.4	14.3	**142.0**	**16.0**	10.2	14.7	39.8	36.1	**2.6**	7.3	52.7	**45.4**
Eaton (ETN.US)	Global giant in electrical management and grid infrastructure; high margins, asset-light, high certainty.	163,036	16.3	10.5	**10.1**	**16.8**	19.8	21.0	23.5	24.5	**2.3**	6.5	31.6	**27.0**
Vertiv (VRT.US)	Absolute leader in liquid-cooling thermal management; dominates the cooling market for NVIDIA (NVDA) Blackwell and other facilities.	125,054	35.6	28.2	**54.0**	**34.0**	23.2	24.5	48.5	46.2	**1.8**	11.6	50.3	**37.6**
Asia Vital Components (3017.TT)	Monopolizes 70% share of AI fans and liquid-cooling key components; tightly tied to NVIDIA (NVDA).	28,337	57.8	35.5	**92.3**	**44.3**	25.0	26.3	58.1	57.1	**4.2**	5.3	24.3	**16.9**
	Simple Average		**50.4**	**44.5**	**293.6**	**60.6**	**27.4**	**31.7**	**30.4**	**33.8**	**1.7**	**15.5**	**56.8**	**36.5**
	Market Cap Weighted Average		**56.9**	**53.4**	**390.2**	**75.9**	**29.7**	**35.6**	**26.8**	**32.0**	**1.4**	**19.3**	**62.6**	**38.5**
(4) Strategic Reshaping														
Intel (INTC.US)	Benefits from US political subsidies for domestic advanced foundry; earnings rebound through the transition.	667,805	11.0	11.5	**157.6**	**43.7**	12.2	15.4	3.6	6.5	**-0.5**	12.7	122.8	**85.4**
Dell (DELL.US)	Co-selling AI servers with NVIDIA (NVDA) ($50bn AI orders in-hand Backlog).	264,563	18.8	51.5	**26.5**	**74.4**	8.7	9.3	-353.7	-422.9	**3.6**	2.1	41.1	**22.8**
	Simple Average		**14.9**	**31.5**	**92.1**	**59.1**	**10.4**	**12.3**	**-175.1**	**-208.2**	**1.6**	**7.4**	**82.0**	**54.1**
	Market Cap Weighted Average		**13.2**	**22.8**	**120.4**	**52.4**	**11.2**	**13.6**	**-97.8**	**-115.4**	**0.7**	**9.7**	**99.6**	**67.7**
(5) China AI														
Tencent (700.HK)	WeChat ecosystem + games = stable cash flow; AI ads monetized with precision; Tencent (00700 HK) Cloud re-rating.	478,767	10.3	9.4	**8.3**	**11.0**	32.6	33.2	19.3	18.6	**6.2**	4.3	11.9	**10.7**
Alibaba (BABA.US/9988.HK)	Alibaba (BABA) Cloud/Qwen benefits from China's AI compute boom; legacy e-commerce as defense.	228,240	2.7	9.6	**-59.0**	**59.4**	6.3	8.1	8.0	8.7	**-3.2**	1.1	20.1	**15.1**
Knowledge Atlas / Zhipu AI	China LLM "new seven little dragons"; flagship domestic generative AI; potential inclusion in the Hang Seng TECH Index.	125,080	346.2	173.0	**-32.8**	**6.2**	-147.0	-55.6	46.3	554.2	**N/A**	1,172.7	N/A	**N/A**
Victory Giant (2476.HK/300476.SZ)	Core leading supplier of high-end PCB/HDI boards for NVIDIA (NVDA) GPUs and AI servers.	46,430	75.1	70.6	**95.8**	**77.6**	31.0	32.3	30.8	38.1	**-1.0**	15.8	33.2	**18.7**
Baidu (BIDU.US/9888.HK)	Cloud revenue +79%; in-house Kunlunxin chip spun off for listing; Apollo Go robotaxi autonomous driving.	35,383	2.7	7.7	**-5.6**	**18.1**	10.2	11.5	5.3	5.9	**-4.1**	1.0	14.0	**11.9**
	Simple Average		**87.4**	**54.1**	**1.3**	**34.4**	**-13.4**	**5.9**	**21.9**	**125.1**	**-0.5**	**239.0**	**19.8**	**14.1**
	Market Cap Weighted Average		**57.4**	**34.9**	**-10.3**	**26.1**	**0.5**	**13.9**	**20.2**	**89.9**	**2.2**	**163.9**	**13.5**	**10.8**
	KGI AI 30 - Simple Average		**62.0**	**41.1**	**169.9**	**45.7**	**27.6**	**33.0**	**26.5**	**39.9**	**1.5**	**50.8**	**45.1**	**31.3**
	KGI AI 30 - Market Cap Weighted Average		**53.0**	**36.3**	**118.2**	**38.6**	**42.4**	**44.3**	**55.4**	**50.9**	**2.0**	**16.6**	**38.2**	**28.1**

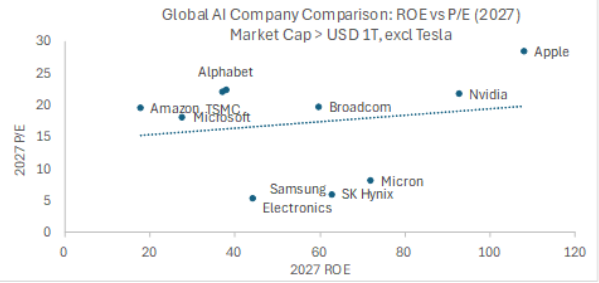

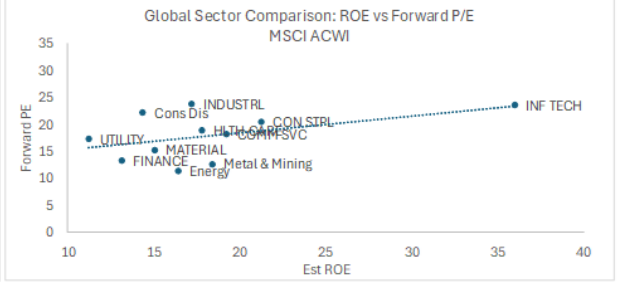

Source: Bloomberg, KGI (As of 26/06/2026)

1) Global AI Infra Monopoly

- **NVIDIA (NVDA):** Monopolizes GPU hardware and global CUDA software development; compute dominance; the majority of industry Capex will translate into NVIDIA's revenue/profit; high margins, reasonable valuation.

- **TSMC (TSM US / 2330 TT):** Monopolizes foundry production of all top-tier AI compute chips globally; possesses world-leading leading-edge wafer fabrication and CoWoS advanced packaging technology; high margins, low valuation.

- **Broadcom (AVGO):** Dominates high-speed networking silicon for high-end AI clusters; exclusive development partner for custom AI chips (ASIC) for Alphabet (GOOGL) and other tech giants; asset-light, high-GM model sustains robust FCF and high dividends.

- **SK Hynix (000660 KS):** No.1 global market share in high-end AI memory (HBM); NVIDIA (NVDA)'s primary memory partner; multi-year order visibility locked in; highest certainty within the AI compute hardware ecosystem outside of NVIDIA (NVDA).

- **Micron (MU):** Leveraging breakthroughs in power-efficient HBM (HBM3e/HBM4), it has entered the top-tier AI compute supply chain, benefiting from a memory supercycle with rising volumes and pricing driven by AI large models' inelastic need for massive, high-bandwidth memory. The only pure-play memory stock in US equities.

- **Samsung Electronics (005930 KS):** Having put in a major bottom after early HBM lag, now at a key inflection point as its next-gen high-end HBM4 memory has cleared qualification with top-tier customers.

- **ASML (ASML):** Monopoly supplier of high-end EUV lithography systems required to manufacture advanced chips—a bottleneck TSMC (TSM US / 2330 TT), Intel (INTC), and Samsung Electronics (005930 KS) cannot bypass when expanding leading-edge capacity; relatively oversold, with earnings estimate revisions lagging.

2) Mega 7

- **NVIDIA (NVDA):** Monopolizes GPU hardware and the global CUDA software ecosystem; compute dominance; the majority of industry Capex will translate into NVIDIA's revenue/profit; high margins; reasonable valuation.

- **Alphabet (GOOGL):** Has demonstrated top-tier capabilities in GenAI through Gemini, effectively monetizing its search engine and massive ad traffic; valuation is relatively inexpensive within the Magnificent Seven.

- **Apple (AAPL):** Owns the world's stickiest hardware-software ecosystem, the natural gateway for AI applications reaching consumers; even if

near-term hardware innovation slows, its robust free cash flow and buyback firepower effectively support share-price performance.

- **Microsoft (MSFT):** Leveraging its deep investment in OpenAI and Azure cloud's dominant edge, it led in seamlessly embedding GenAI into Office and other enterprise software, underpinning strong B2B pricing power and compounding growth.

- **Amazon (AMZN):** Its core profit engine, AWS cloud computing, has seen a strong rebound in growth, while the e-commerce business has delivered significant margin expansion after completing supply-side optimization—making it a balanced recovery play.

- **Meta Platforms (META):** AI is significantly enhancing ad targeting precision and monetization on its FB/IG social platforms, while acquisitions of new technologies (e.g., Manus AI) continue to broaden its hardware and software boundaries.

- **Tesla (TSLA):** Even as pure EVs face cyclical headwinds, its core value proposition has fully pivoted to Full Self-Driving (FSD), energy storage, and humanoid robots—making it the highest-beta future-tech stock among the Magnificent Seven.

3) Segment Leaders

- **AMD (AMD):** NVIDIA (NVDA)'s direct competitor, staying close with the MI300/400 series; the bet is on capturing spillover compute orders from a nonlinear inflection and achieving higher share-price beta.

- **Renesas (6723 JP):** Holds formidable moats in the automotive supply chain; amid a JPY asset re-rating, a slow-bull play with dual exposure to Edge AI and intelligent-vehicle cycles, offering solid downside protection.

- **SanDisk/Western Digital (SanDisk - SNDK US):** A classic "cycle upturn + AI must-have" high-beta stock; explosive data center demand for high-end, high-capacity SSDs is propelling it into an earnings boom phase.

- **Marvell (MRVL):** In the high-speed optical module/DSP market (within AI servers, electrical signals are converted to optical for transmission and managed via DSP chips for control and error correction), Marvell and Broadcom (AVGO) form a duopoly. It is also a core partner for Big Tech's in-house ASIC AI chips. High GM, asset-light.

- **Coherent (COHR):** One of the world's largest suppliers of optical modules and optoelectronic components, integrating optical chips, housings, and interfaces into complete optical modules (Transceivers) sold to data centers. It is also a global giant in wafer materials such as SiC and Indium Phosphide (InP), controlling the most upstream materials processes for optoelectronic conversion. Asset-intensive with heavy depreciation; once 1.6T optical modules reach mass production, net profit should surge nonlinearly.

- **Lumentum (LITE):** Monopolizes global EMI physical photonics chips—the core, highest-barrier light-emitting component in optical modules. Without high-end optical chips, optical modules cannot emit optical signals. As line rates rise from 800G to 1.6T (a 2x step), demand for high-end EML chips and CW (continuous-wave) high-power lasers increases inelastically, roughly doubling. Its de facto monopoly ties order visibility directly to the global compute capacity buildout.

- **GE Vernova (GEV US):** An independently listed pure-play power heavy equipment name spun off from GE; owns the world's most advanced, highest-efficiency HA-class gas turbines (sub-one-minute fast start). It helps Hyperscalers avoid waiting 3–5 years for traditional grid expansions by directly procuring utility-scale generator sets to bring data centers online faster; a high-visibility energy name (backlog extends several years out).

- **Eaton (ETN):** Focused on medium- and low-voltage power distribution and in–data center electrical architecture (transformers stepping high voltage down to equipment-ready power, UPS, and intelligent rack PDU). A prime play on the AI energy bottleneck: high-margin, asset-light (no heavy generator-set manufacturing; focused on last-mile distribution hardware near servers and power management software); high visibility (US data center expansions have upgrade orders for distribution systems and transformers booked out for years).

- **Vertiv (VRT):** The absolute leader in high-temperature liquid-cooling thermal management, benefiting as high-power AI chips from NVIDIA (NVDA) and others trigger massive, inelastic demand for data-center liquid cooling; high fundamental visibility (independent of competition among LLMs—so long as compute capacity expands, customers must order the company's products).

- **AVC (3017 TT):** Leveraging a dual-track strategy of 3D VC and cold plates to become the thermal leader among Taiwan-listed stocks; tightly aligned with NVIDIA (NVDA)'s latest server architecture upgrade cadence, enabling faster technology deployment than US-listed peers.

4) Strategic Reshaping

- **Intel (INTC):** In a major transformation, benefiting from policy-driven US subsidies for domestic advanced foundry and from a commercial "AI PC" refresh cycle that offers strong margin recovery leverage; earnings are poised to rebound.

- **Dell Technologies (DELL):** The go-to hardware partner for traditional government and enterprise customers to operationalize AI, leveraging a deep partnership with NVIDIA (NVDA) to sell AI servers. In a high-valuation tech environment, its robust, stable cash flows, planned share buybacks, and a 27% payout ratio, together with a reasonable valuation range, offer investors attractive risk-reward.

5) China AI

- **Tencent (00700 HK):** Leveraging WeChat's uniquely sticky ecosystem and massive cash flows, accelerating two-sided monetization across B/C via overseas gaming expansion and the Hunyuan LLM; a value stock with strong anti-cyclical resilience and AI-driven incremental upside, trading at an attractive valuation.

- **Alibaba (BABA / 09988 HK):** Core e-commerce bottoming out with profits and market share recovering; rapid overseas cross-border e-commerce expansion; Alibaba Cloud set to fully benefit from the surge in China enterprise AI compute; low valuation and ongoing share buybacks provide a high margin of safety; valuation remains cheap.

- **Baidu (BIDU / 09888 HK):** China's GenAI commercialization leader, reshaping search advertising and cloud via LLMs; autonomous driving has reached scaled commercial operations across multiple cities nationwide; a high-moat tech stock with strong valuation optionality, oversold and cheap.

- **Victory Giant Tech (02476 HK / 300476 SZ):** As the leading indirect PCB supplier to NVIDIA (NVDA) for core AI servers and GPUs, it is a direct beneficiary of the global AI compute arms race; surging orders for high-layer-count, high-density PCBs are driving results; dual listings in Hong Kong and Mainland China support valuation re-rating potential.

- **Zhipu AI (2513 HK):** The Tsinghua-lineage player with the deepest technical roots among China's "New Seven Little Dragons" of LLMs; GLM full-stack, self-developed foundation model; dual backing from the national team and Big Tech; B2B commercialization across all scenarios; exceptional ecosystem penetration across academia, government/enterprise, and consumer. A China GenAI bellwether with the highest technical certainty and the most robust closed-loop commercialization.

1. Overview of Global AI Ecosystem (20 Core Layers)

I. Upstream						
#	**Global AI tiers**	**Importance of this level**	**Key sub-sectors**	**Bottleneck Status**	**Challenges / Current Bottlenecks**	**Major Companies and Their Businesses**
1	**Energy and Power Supply**	AI development is constrained by power supply; training frontier models and operating hyperscale data centers require GW-scale power. Without stable, scalable, and increasingly clean energy, AI's physical infrastructure cannot continue to expand.	A) Utility-scale power generation (grid-level) B) Renewable Energy and Nuclear C) Power Stability and Redundant Backup D) Power Purchase Agreements (PPA) and Energy Trading	**Current bottlenecks are hampering**	AI data centers' power demand is growing far faster than grid capacity expansion. Approvals and construction for new generation and transmission infrastructure take 5–10 years, the toughest physical bottleneck to AI's at-scale expansion today.	**GE Vernova** (GEV US) / **Siemens Energy** (ENR GY) / **Constellation Energy** (CEG US) / **Vistra** (VST US) / **Schneider Electric** (SU FP) / **NextEra Energy** (NEE US)
2	**Advanced Materials and Chemicals**	Semiconductor manufacturing requires the highest-purity materials on earth. Advanced AI chips are entirely dependent on specialty silicon wafers, photoresists, and rare gases. If the supply of these foundational materials is disrupted, chip production will halt immediately.	A) Silicon Wafers B) Photoresist C) Specialty gases D) Rare Earth Elements E) Advanced Substrates	**No bottlenecks for now**	Not a bottleneck for now, but supply is highly concentrated. High-purity chemicals are sourced primarily from Japan; any disruption would quickly bring global chip production to a halt.	**Shin-Etsu Chemical (4063 JP) / Entegris (ENTG US) / SUMCO (3436 JP) / JSR** (4185 JP) / **Air Liquide / Linde** (AI FP) / **Tokyo Ohka Kogyo (4186 JP) /**
II. Semiconductor Manufacturing Backbone						
#	**Global AI Tiers**	**Importance of this level**	**Key sub-sectors**	**Bottleneck status**	**Challenges / Current Bottlenecks**	**Key companies and their operations**
3	**Semiconductor manufacturing equipment**	These companies manufacture the machines that print microscopic circuits onto silicon wafers—an ultimate strategic chokepoint. Without ASML (ASML)'s extreme ultraviolet (EUV) lithography systems, no one in the world can produce cutting-edge AI chips.	A) Lithography equipment (EUV/DUV) B) Thin-Film Deposition Equipment C) Etching equipment D) Metrology and Inspection Equipment	**Gradually easing but still tight**	ASML (ASML) is the sole supplier of EUV tools. These systems, costing over USD 300 million each, have extremely long lead times and are subject to geopolitical export controls, creating a persistent strategic bottleneck.	**ASML — Sole-source EUV supplier (ASML) / Applied Materials** (AMAT US) / **Lam Research** (LRCX US) **Tokyo Electron (8035 JP) / Advantest (6857 JP) / KLA (KLAC US)**
4	**Semiconductor Wafer Fabs / Foundries**	Wafer fabs, factories that cost USD tens of billions and turn chip blueprints designed by companies like NVIDIA (NVDA) into physical silicon wafers, are among the most capital-intensive manufacturing processes in human history.	A) Advanced Logic Process B) Advanced Packaging C) Yield optimization	**Once the key bottleneck**	Building an advanced wafer fab costs over USD 20 billion and takes several years. Supply is highly concentrated in Taiwan, leaving the global AI supply chain highly vulnerable to geopolitical tensions or natural disasters.	**TSMC (TSM US) / Samsung Foundry (005930 KS) Intel Foundry (INTC US) / SMIC (981 HK) GlobalFoundries** (GFS US)
5	**Advanced packaging and interconnect**	With transistor scaling slowing, the industry is shifting to tightly integrating multiple small dies (chiplets). Advanced packaging ensures inter-die communication is fast enough for the system to perform like a powerful AI brain.	A) CoWoS / HBM stacked packaging B) Chiplets C) Interposers D) OSAT — Outsourced Assembly and Test	**Current bottlenecks are hindering**	TSMC's (TSM US / 2330 TT) CoWoS packaging capacity was the primary bottleneck for NVIDIA GPU supply in 2023–2024. Given the complex packaging flow and low yields, even ongoing capacity expansions at TSMC continued to constrain global chip shipments.	**TSMC (TSM US) — CoWoS packaging / ASE Technology (ASX US) / Amkor (AMKR US) / Samsung Advanced Packaging** (005930 KS) / **JCET Group (600584 CH)**
III. Core chip components (inside AI chips)						
#	**Global AI Stack**	**Importance of this level**	**Key sub-sectors**	**Bottleneck status**	**Challenges / Current Bottlenecks**	**Major companies and their businesses**
6	**Logic chips**	GPU and custom accelerators are the true "brain" of AI, purpose-built to perform the massively parallel mathematical computations required to train neural networks and generate responses.	A) Graphics Processing Unit (GPU) B) ASIC / Custom Accelerators C) Central Processing Unit (CPU) D) Dedicated inference chips	**Once the key bottleneck**	Demand for NVIDIA (NVDA)'s latest architecture continues to outstrip supply. The core challenge is managing the extreme heat and power consumption of next-generation chips while scaling to hyperscale clusters.	**NVIDIA** (NVDA US) / **AMD** (AMD US) / **Intel** (INTC US) / **Google TPU** (GOOGL US) / **Broadcom ASIC** (AVGO US) / **Baidu (BIDU / 09888 HK) — Kunlun (Baidu Kunlun chip) / Apple — AI SoC (AAPL US) / Renesas (6723 JP)**
7	**Memory**	AI processors compute at extremely high speeds but require a continuous stream of input data. High Bandwidth Memory (HBM) is installed adjacent to the processor to move data at very high speed. If memory capacity is insufficient, costly GPUs will sit idle waiting for data.	A) High Bandwidth Memory (HBM) B) Dynamic Random-Access Memory (DRAM) C) NAND Flash	**Gradually easing But still tight**	HBM is difficult to manufacture and requires advanced packaging to interface with GPUs. Supply of HBM3 and HBM3e is severely constrained, directly capping global AI accelerator shipment volumes.	**SK Hynix — HBM leader (000660 KS) / Micron (MU US) Samsung Electronics** (005930 KS) / **SanDisk** (SNDK US)
8	**AI Cluster Networking and Interconnects**	A single chip cannot train modern AI models. Tens of thousands of GPUs must be interconnected to operate as one supercomputer. A high-speed network enables real-time data sharing across chips, minimizing latency during training.	A) High-speed Ethernet B) Optical interconnect C) NVLink-class chip interconnects D) Printed Circuit Board (PCB)	**Current bottlenecks are impeding**	As cluster sizes expand beyond 100,000 GPUs, network interconnect becomes the bottleneck. Electrical links generate excessive heat and latency, forcing a difficult industry transition to optical interconnects (silicon photonics).	**Broadcom** (AVGO US) / **NVIDIA — NVLink / InfiniBand** (NVDA US) / **Marvell** (MRVL US) / **Lumentum** (LITE US) / **Amphenol** (APH US) / **Zhongji Innolight** (300308 CH) / **Gold Circuit Electronics (2368 TT) / Victory Giant Technology (300476 CH) / Unimicron (3037 TT) / Coherent** (COHR US) / **Cisco** (CSCO US) / **Arista Networks** (ANET US)
IV. Data Center Physical Infrastructure (AI's Habitat)						
#	**Global AI tiers**	**Importance of this level**	**Key sub-segments**	**Bottleneck status**	**Challenges / Current Bottlenecks**	**Major Companies and Their Businesses**
9	**Power T&D and Electrical Equipment**	Grid power cannot be fed directly into servers. It must be stepped down, conditioned, and distributed to each rack with extremely high stability. Power equipment ensures large amounts of electricity are safely delivered to server racks, preventing catastrophic outages.	A) Uninterruptible Power Supply (UPS) B) Transformers C) Power Distribution Unit (PDU)	**Current bottlenecks are hindering**	Lead times for large power transformers and specialized switchgear have stretched from months to years. Without these, data centers cannot be energized, making them a severe supply-chain bottleneck.	**Schneider Electric (SU FP) / Eaton (ETN US) / ABB (ABBNY US) / Vertiv (VRT US) / Delta Electronics (2308 TT) / GE Vernova — Grid Solutions (GEV US)**
10	**Cooling system**	AI chips generate massive heat. Conventional air cooling can no longer handle the latest high-density GPU racks. Advanced liquid-cooling is now essential to prevent chip overheating and ensure safe data center operations.	A) Liquid Cooling (Direct-Contact) B) Immersion cooling C) Thermal management	**Current bottlenecks Are hindering**	Transitioning from air cooling to liquid cooling requires a comprehensive retrofit of data-center piping and rack designs. Manufacturing capacity for cooling distribution units (CDUs) is struggling to keep pace with NVIDIA (NVDA) Blackwell deployments.	**Vertiv (VRT US) / Asia Vital Components (3017 TT) / Auras Technology (3324 TT) / Asetek (ASETEK NO) / Supermicro (SMCI US)**
11	**Data Center Build-out and Operations**	These are AI's physical fortresses, providing highly secure, structurally reinforced, high-power-density environments to house tens of thousands of heavy, high-power server racks.	A) Wholesale and Retail Colocation (Colocation) B) Real Estate Investment Trusts (REITs) C) Sovereign Data Centers D) AI Server Build-out and Deployment	**Current bottlenecks are hindering**	Finding sites with both large-scale power interconnection and high-speed fiber is increasingly difficult. Build-out is constrained by zoning regulations, grid interconnection queues, and equipment shortages across multiple layers.	**Equinix** (EQIX US) / **Digital Realty** (DLR US) / **Dell** (DELL US) **HP Enterprise (HPE US) / Supermicro (SMCI US) / Hon Hai (2317 TT) / NTT Data (9613 JP) / Wiwynn (6669 TT) / Quanta Computer (2382 TT) / GDS (GDS US)**

V. AI Infrastructure and Cloud Platforms						
#	**Global AI Stack**	**Importance of this Level**	**Key sub-sectors**	**Bottleneck status**	**Challenges / Current Bottlenecks**	**Key companies and their businesses**
12	**Hyperscale Cloud Service Providers and AI Cloud Platforms**	Tech giants rent high-performance GPU clusters on demand; they are the primary buyers of AI accelerators and the backbone of AI infrastructure. As gatekeepers of compute, they enable developers and enterprises to access supercomputing-class capacity without building their own data centers.	A) Infrastructure as a Service (IaaS) B) Platform as a Service (PaaS) / AI Cloud C) Sovereign / Compliance Cloud Services D) Managed AI model APIs (e.g., AWS Bedrock, Azure OpenAI)	**No bottlenecks for now.**	While raw compute is increasingly abundant, high costs and vendor lock-in remain challenges. Enterprises face high GPU rental costs, and competition among hyperscalers is intensifying as each seeks to lock customers into its own ecosystem.	**Microsoft Azure** (MSFT US) / **Amazon AWS** (AMZN US) / **Google Cloud** (GOOGL US) / **Oracle Cloud** (ORCL US) / **Alibaba Cloud (BABA US) / Baidu Cloud (9888 HK)**
13	**GPU Cloud and Professional Inference Platforms**	Pure-play GPU cloud providers and specialized inference platforms offer compute capacity rental tailored to AI workloads. Unlike diversified hyperscalers, these firms focus on speed, cost, and AI-optimized infrastructure, serving labs and enterprises that need massive raw compute or fast model inference.	A) GPU-as-a-Service (pure GPU cloud) B) Managed inference endpoints C) Local GPU Clusters D) Inference acceleration platforms, such as Groq LPU	**No bottlenecks for now.**	As hyperscalers expand their own GPU cloud offerings, competition is intensifying. Pure-play GPU cloud providers must differentiate on price, provisioning speed, and AI-specific optimization to avoid commoditization.	**CoreWeave Cloud (CRWV US) / Nebius (NBIS US) / Groq API (Private) / OVHcloud (OVH FP) — European pure GPU cloud**
14	**Cybersecurity and AI Security**	As AI becomes deeply embedded in critical infrastructure, it becomes a high-value attack target. This layer protects the entire tech stack—from the network perimeter and cloud workloads to the model itself—against hacking, data theft/exfiltration, adversarial attacks, and prompt injection.	A) Endpoint and Cloud Workload Protection B) Network and Perimeter Security C) Identity and Access Management D) Data Loss Prevention	**Future bottlenecks**	AI systems introduce a new attack surface—adversarial inputs, model inversion, and supply chain poisoning—threats traditional cybersecurity tools weren't designed to address. The security industry is scrambling to catch up.	**CrowdStrike** (CRWD US) / **Palo Alto Networks** (PANW US) **Cloudflare** (NET US) / **Fortinet** (FTNT US) / **Microsoft Defender** (MSFT US)
VI. Models, Software, and Data (Visible AI Layer)						
#	**Global AI hierarchy**	**Importance of this level**	**Key sub-sectors**	**Bottleneck Status**	**Challenges / Current Bottlenecks**	**Major companies and their businesses**
15	**Foundation models**	Neural networks with billions of parameters trained on massive datasets exhibit broad reasoning capabilities. This layer is dominated by capital-intensive competition, with an ongoing contest between closed-source proprietary models and open-weights community development.	A) Proprietary / Closed-Source Models B) Open Weights / Open-Source Models C) Verticals / Edge Models	**No current bottlenecks.**	The "data wall" is looming—high-quality human text for training is running out. Training costs have surged into the USD billions, meaning only a handful of deep-pocketed institutions can compete at the frontier.	**OpenAI ChatGPT (Private) / Google Gemini (GOOGL US) / Anthropic Claude (Private) / Meta Llama (META US) / MiniMax Hailuo (100 HK) / DeepSeek (Private) / Moonshot AI — Kimi (Private) / Zhipu AI (2513 HK) / xAI — Grok (Private)**
16	**AI Frameworks and Tools**	MLOps libraries and orchestration tools turn brittle raw models into reliable production applications, abstracting complex mathematical operations into repeatable, deployable code.	A) Core Machine Learning Frameworks B) Orchestration and Intelligent Agent Framework C) Model Repositories and Model Marketplaces D) MLOps and Evaluation Tools	**No bottleneck for now**	The ecosystem is highly fragmented and evolving too fast. Developers face tool fatigue and struggle to build reliable production systems as underlying models and frameworks churn every few months.	**Google TensorFlow** (GOOGL US) / **Hugging Face** (Private) / **Databricks — MLflow** (Private) / **LangChain** (Private) **AutoGen — Microsoft** (MSFT US) **Weights & Biases (W&B) (Private) — Leading MLOps platform**
17	**Data Platforms and Synthetic Data**	Collect, organize, and clean enterprise datasets to build a unified data stream. A model's intelligence hinges on the quality of its training data—low-quality inputs will trigger hallucinations and reasoning errors.	A) Data Lakehouse / Cloud Data Warehouse B) Data labeling and curation C) Vector Databases and Knowledge Storage D) Synthetic data generation	**Future bottlenecks Compute and models will become abundant, but data remains messy**	Enterprise data is messy, siloed, and unstructured. The core challenge is securely connecting proprietary enterprise data to AI models without exposing sensitive information, while continuously sourcing high-quality data to drive sustained intelligence improvement.	**Snowflake** (SNOW US) / **Palantir** (PLTR US) / **Databricks** (Private) / **Scale AI** (Private) / **SAP** (SAP US)
VII. Downstream — Applications and Commercialization						
#	**Global AI tiers**	**Importance of this tier**	**Key sub-sectors**	**Bottlenecks**	**Challenges / Current Bottlenecks**	**Key companies and their businesses**
18	**Horizontal Enterprise AI**	Cross-industry productivity tools for the GenAI era, optimizing common workflows such as document drafting, coding, and summarization, reclaiming millions of work hours for enterprises.	A) AI Co-pilot and Assistant B) Autonomous Workflow Automation C) Creative & Content Suite	**No bottlenecks for now.**	The primary challenges are user adoption rates and proving tangible ROI. Despite powerful tools, integrating them into legacy enterprise workflows and daily life—and changing habits—remains a major commercial hurdle.	**Microsoft — M365 Copilot (MSFT US) / Salesforce — Einstein (CRM US) / Adobe — Firefly (ADBE US) / ServiceNow (NOW US) / Tencent — WeChat AI (700 HK) / Apple — Apple Intelligence (AAPL US) / Intuit (INTU US) / Manus AI (Private)**
19	**Vertical / Industry AI**	Highly specialized AI for high-risk domains such as healthcare and finance requires deterministic accuracy and deep integration with regulated data ecosystems—areas where general-purpose models fall short.	A) Healthcare and Biopharma B) Embodied AI, Mobility, and Automation C) Financial & Compliance Intelligence D) Legal and Professional Services AI	**Future bottlenecks Regulatory approval queue Continuing to lengthen**	Regulated industries demand deterministic accuracy and explainability. The bottlenecks are lengthy regulatory approval backlogs (e.g., FDA approvals for AI medical devices) and limited access to highly protected, proprietary industry data.	**Healthcare: Tempus (TEM US) / Siemens Healthineers (SHL GY) Financials: Bloomberg (Private) / Visa (V US) Manufacturing: Siemens (SIE GY) / Rockwell (ROK US) Robotics: Ubtech Robotics (9880 HK) Mobility: Waymo (GOOGL US) Tesla (TSLA US) Legal: Harvey AI (Private)**
VIII. Governance, Security, and Regulation						
#	**Global AI Tiers**	**Importance of This Level**	**Key Sub-Sectors**	**Bottleneck Status**	**Challenges / Current Bottlenecks**	**Example**
20	**Governance, Security and Regulation**	Establish global guardrails across legal and software layers to prevent bias, misinformation, and intellectual property theft, ensuring AI models are safe, transparent, and compliant before deployment.	A) Corporate Audit and System Integration B) Alignment Research and Threat Mitigation C) Regulatory Compliance and Legal Framework D) Export Controls and Geopolitical Governance	**Future bottlenecks**	Regulatory progress severely lags the pace of AI development. Overregulation could stifle innovation; underregulation risks catastrophic security breaches, IP theft, and deployment of unsafe autonomous systems.	**EU AI Act (European Union AI Act) US NIST AI RMF (Risk Management Framework, risk management program)**

Source: KGI

2. The Four Strategic Pillars

In the global AI ecosystem, Power sets the capacity ceiling, Foundry determines compute supply, Networking/Storage sets computational efficiency, and LLM/Agent determines ultimate monetization and profit capture. In asset allocation, leaders across these four sub-segments exhibit exceptional pricing power and resilience.

Global AI Ecosystem: Panoramic View of the Four Strategic Pillars

Strategic segment	Strategic Importance	Underlying Logic	Core Companies
1. Energy & Nuclear / Power Infrastructure	Power or no power? The absolute survival threshold for compute build-outs; power availability directly determines whether data centers can be built.	Hyperscale data centers consume as much power as mid-sized cities. Zero-carbon, 24/7 reliable nuclear power and Small Modular Reactors (SMR) are scarce. Tech giants are signing multi-decade contracts to pre-emptively lock in grid capacity.	Constellation Energy (CEG), GE Vernova (GEV), NextEra Energy (NEE)
2. Foundry and Advanced Packaging	Chip or No Chip? The Ultimate Moat in the Global Tech Supply Chain; Can Chips Be Manufactured?	Virtually 100% of top-tier AI accelerator chips—whether from NVIDIA (NVDA), AMD (AMD), or Big Tech ASICs—rely exclusively on TSMC (TSM US / 2330 TT) for advanced nodes and CoWoS packaging. If geopolitics or capacity bottlenecks flare up, the evolution of global AI compute would instantly come to a physical standstill.	TSMC (TSM US / 2330 TT)
3. High-Bandwidth Memory and High-Speed Networking	I/O speed: It determines whether chips can run at high speed. The most severe bottleneck today.	Without HBM (high-bandwidth memory), a GPU is a data-starved shell; without high-speed interconnect chips and switching fabric, tens of thousands of chips can't be fused into a supercomputer.	SK Hynix (000660 KS), Micron (MU), Broadcom (AVGO)
4. Core LLMs and Agent Ecosystem	Applications? The ultimate traffic gateway and ecosystem moat;	Frontier LLMs are rapidly evolving into Multi-Agent Systems (MAS), directly taking over enterprise productivity workflows (e.g., charging by Token throughput). This will set the monetization ceiling for the entire AI industry. Once B2B enterprises integrate deeply, switching costs become prohibitive, effectively handing over the control plane for workflows and financial data.	Anthropic, OpenAI, Microsoft (MSFT) (Copilot)

Source: KGI

2.1 AI Clean Energy & Power Grid

AI ultimately reduces to compute, and compute ultimately reduces to power. A hyperscale AI data center with 100,000 to even 1,000,000 top-tier chips can consume as much electricity as a medium-sized city. Whoever can provide stable, climate-independent, zero-carbon 24/7 baseload power (especially nuclear and SMR) effectively holds the approval power over whether AI data centers get built. That's why Hyperscalers like Microsoft (MSFT) and Amazon (AMZN) are signing multi-decade contracts to lock in nuclear capacity. Energy is not just a cost line—it is the hard survival constraint for scaling AI compute.

Key players: Constellation Energy (CEG), GE Vernova (GEV), NextEra Energy (NEE)

2.2 Advanced Foundry & Packaging

This sub-industry has the highest barriers to entry and the strongest bargaining power in the global technology supply chain. Whether it is NVIDIA (NVDA)'s cutting-edge Rubin/Blackwell chips, AMD (AMD)'s MI series, or custom AI chips (ASIC) from Alphabet (GOOGL) (TPU) and Amazon (AMZN) (Trainium), virtually 100% of top-tier AI accelerator chips globally must rely on TSMC (TSM US / 2330 TT)'s advanced process nodes (3nm / 2nm) and CoWoS advanced packaging technologies to be fabricated.

Software algorithms can be refactored in code at any time, but if geopolitics or capacity bottlenecks trigger foundry shutdowns, the evolution of global AI compute will abruptly hit a physical standstill.

Core players: TSMC (TSM US / 2330 TT)

2.3 HBM & High-Speed Networking Fabric

This is a critical lever for breaking the physical bottlenecks to compute. Today, the most severe constraints on large-model inference and large-scale training are often not insufficient GPU compute, but lagging data-transfer bandwidth—both on-chip and across hyperscale server clusters—i.e., the Memory Wall and the Network Wall.

In the current HBM4 (high-bandwidth memory) era, without SK Hynix (000660 KS) or Micron (MU) supplying ultra-high-bandwidth memory chips, GPUs are compute-rich yet data-starved; likewise, without Broadcom (AVGO) providing customized interconnect silicon and ultra-low-latency switching architectures, tens of thousands of GPUs cannot be fused into a true supercomputer. These are the "uncrowned kings" hidden behind the headline hardware.

Core players: SK Hynix (000660 KS), Micron (MU), Broadcom (AVGO)

2.4 Frontier LLM & Multiagent Platforms

Frontier LLMs are the central operating system of the information age and the ultimate traffic gateway at the business layer. This sub-sector determines the AI industry's monetization ceiling and data control rights. Gartner forecasts indicate that as LLMs rapidly shift toward Multiagent Systems (MAS) and enterprise-grade domains (e.g., Anthropic, via Claude Code, is directly taking over coding workflows at global large enterprises and charging based on Token throughput), whoever controls the base-layer inference model will be the gatekeeper to global enterprise workflows and financial data over the next decade. Once B2B enterprises are deeply integrated with a given model, Switching Costs are extremely high, enabling model giants to build massive ecosystem moats.

Key players: Anthropic, OpenAI, Microsoft (MSFT) (Copilot)

3. The Top-Five Global Leaders

In the global AI ecosystem, the 2026 market structure has evolved into a highly interdependent and highly co-cooperative empire. **ASML (ASML) (exclusive supplier of EUV lithography systems) -> TSMC (TSM US / 2330 TT) (exclusive foundry and advanced packaging) -> NVIDIA (NVDA) (compute leader) / Broadcom (AVGO) (networking leader) / AMD (AMD) (NVIDIA rival)**

The 5 undisputed global AI infrastructure giants (G5)

The Big Five	Industry Roles	Core Moat (Moat)	Core business	Profitability
NVIDIA (NVDA)	Fabless (system vendors); the absolute dominant player in today's compute market	Irreplaceable GPU hardware assets and the CUDA software ecosystem.	Full-stack AI computing platform (chips, NVLink network, server systems, CUDA software)	The king of windfall profits, gross margin as high as 75%.
TSMC (TSM US / 2330 TT)	Foundry; the most defensive core assets in the compute supply chain	Virtually monopolizes over 99% of global leading-edge AI chip foundry and packaging, serving NVIDIA (NVDA), AMD (AMD), Apple (AAPL), etc.	Leading-edge logic chip manufacturing (3nm/2nm) and advanced packaging (CoWoS, SoIC)	Exceptionally robust profitability profile; GM at 66.2%.
ASML	Hardware (semiconductor equipment); benefiting from capacity bottlenecks, order backlog already fully booked in advance.	Sole supplier of the ultimate strategic equipment (EUV lithography systems) essential for manufacturing advanced AI chips.	Lithography systems (EUV, High-NA EUV)	GM ~50%; NM ~30%.
Broadcom (AVGO)	Fabless (customized/network); core beneficiary of network architecture upgrades	Core Ethernet switching chips that absolutely dominate 100k-card-class hyperscale data centers.	Custom AI ASICs (e.g., Alphabet (GOOGL) TPU, Meta (META) accelerators) and high-bandwidth network switch chips (Tomahawk series)	Exceptionally strong profitability; GM as high as 65%–70%.
AMD	Fabless (system vendors); Highly flexible potential challengers	The industry's only compelling price/performance alternative capable of going head-to-head with NVIDIA (NVDA)'s flagship GPU.	Compute alternatives (MI300/MI400 series accelerators) and x86 server CPUs	Current margins below peers, but expected to expand.

Source: KGI

3.1 The Up/Downstream Ties and Competition

The five giants (G5) are far from isolated competitors; they form an interlocking, highly symbiotic yet mutually checking "ecosystem food chain." They feature both absolute upstream/downstream dependencies and highly nuanced coopetition. Below, we analyze their business and dependency across three core dimensions:

First: Core physical supply chain—an absolute upstream/downstream dependency chain. ASML (ASML) (sole supplier of EUV lithography) —>

TSMC (TSM US / 2330 TT) (exclusive foundry and advanced packaging) —> NVIDIA (NVDA) / AMD (AMD) / Broadcom (AVGO)

1. Equipment monopoly: ASML (sole supplier of EUV lithography) is TSMC's "mother of industry." To mass-produce 3nm/2nm chips at scale in 2026 and develop next-gen nodes, TSMC must purchase ASML's monopolized High-NA EUV (high-numerical-aperture lithography).

Dependence: 100%. If ASML cuts supply or capacity falls short, TSMC's node advancement would immediately stall.

2. Ultimate gatekeeper of capacity: TSMC (exclusive foundry and advanced packaging) is the primary supplier to NVIDIA, AMD, and Broadcom (three Fabless design houses). NVIDIA's Blackwell/Rubin architecture GPUs, AMD's MI300/MI400 accelerators, and Broadcom's custom chips for Alphabet (GOOGL) TPU and Meta (META) ASICs all rely on TSMC's advanced nodes (4nm/3nm) for foundry.

More critically, in 2026 the bottleneck for AI chips is advanced packaging (CoWoS). NVIDIA, AMD, and Broadcom are vying for TSMC's very limited CoWoS capacity.

Dependence: Extremely high. Although these three design houses are formidable in software and architecture, without TSMC allocating Wafer and packaging capacity, they can't ship a single chip.

Second: Compute and networking—NVIDIA (compute leader) and Broadcom (networking leader) are highly complementary yet fiercely competitive.

1. Business complementarity (jointly underpinning 10k-card clusters). When tens of thousands—even hundreds of thousands—of chips are linked into a giant AI cluster in a non-NVIDIA full-stack AI data center, NVIDIA provides the compute (GPU die), while Broadcom provides backbone networking (e.g., Tomahawk 5 switch silicon, PCIe switch chips, and optical module core components).

Their chips must interoperate perfectly on the server motherboard for data to flow between chips at ultra-low latency.

2. Strategic confrontation (battle of technical roadmaps). NVIDIA aims for a "fully closed loop": it aggressively promotes its NVLink architecture and InfiniBand networking, attempting to bundle compute and networking into a "black box" for customers and crowd Broadcom out.

Broadcom leads the "open counterattack": Broadcom has teamed up with AMD, Alphabet (GOOGL), Meta (META), and others to form UALink and the Ultra Ethernet Consortium, pushing Ethernet to replace NVIDIA's networking and to decouple NVIDIA's lock-in.

Third: Direct competition and coopetition—AMD and Broadcom compete directly with NVIDIA; TSMC is a neutral supplier.

1. NVIDIA (NVDA) vs AMD (AMD): a pure zero-sum rivalry. AMD's MI400 accelerators are closing in on NVIDIA's Blackwell in hardware specs and price-performance. They compete head-to-head for procurement budgets at hyperscalers (Microsoft (MSFT), Oracle (ORCL), etc.).

2. NVIDIA (NVDA) vs Broadcom (AVGO): Broadcom doesn't sell general-purpose GPUs directly, but it is the technology manufacturing partner and co-designer for Alphabet (GOOGL) TPU and Meta (META) AI chips (Custom ASIC business). Alphabet and Meta don't want to be entirely at NVIDIA's mercy, so they engage Broadcom to develop in-house chips. For every additional TPU Broadcom builds for Alphabet, that's one fewer GPU NVIDIA sells.

3. TSMC (TSM US / 2330 TT)'s "absolute neutrality": whether NVIDIA (NVDA) (the ruler), AMD (AMD) (the challenger), or Broadcom (AVGO) (the behind-the-scenes hand helping customers self-develop), they're all TSMC's VIP clients. TSMC favors no side—no matter how you fight, whoever wins still comes here and pays the foundry bill. TSMC will even deliberately support AMD and Broadcom to prevent NVIDIA from becoming too dominant and squeezing TSMC's foundry pricing.

3.2 The Key Success Factors

- **NVIDIA (NVDA): the CUDA ecosystem's ironclad moat.** NVIDIA is more than a chip company. Its CUDA software ecosystem invested in 15 years ago locks millions of AI developers worldwide into running code only on NVIDIA hardware. By 2026, its ability to stitch tens of thousands of chips into "one supercomputer" via NVLink networking makes its cluster efficiency unmatched.

- **TSMC (TSM US / 2330 TT): impeccable commercial credit and the "king of packaging."** TSMC sticks to its promise of "not competing with customers," earning global trust. More importantly, as Moore's Law slows, AI performance hinges on packaging; with CoWoS and SoIC advanced packaging, TSMC holds the unique know-how to assemble GPUs with HBM memory.

- **ASML (ASML): an irreplaceable feat of physics.** ASML's success lies in absolute dominance of extreme ultraviolet (EUV) lithography. Making AI chips is like engraving a space shuttle on a strand of hair; only ASML can produce this level of "divine light" precision.

- **Broadcom (AVGO): network monopoly and custom silicon.** As AI clusters scale from tens of thousands of cards toward hundreds of thousands/millions, the network's value can exceed the chips themselves. Broadcom leverages absolute dominance in Ethernet switch silicon and contract-manufactures custom ASICs for Google and Meta, becoming the low-key cash machine of AI infrastructure.

- **AMD (AMD): the "best Plan B" for a world long suffering under green.** AMD's success stems from superior hardware design (e.g., pioneering advanced Chiplet in the MI300 series) and a "No.2 supplier" strategic positioning. Any hyperscaler unwilling to let NVIDIA choke them will throw orders to AMD.

3.3 Profitability Comparison

- **NVIDIA (NVDA) (king of windfall profits):** GM stays at 75%–80%, NM exceeds 55%. As a systems integrator, it sells an all-in bundle spanning software, networking, and compute, wielding exceptional monopoly pricing power.
- **TSMC (TSM US / 2330 TT) (the solid shield):** GM around 53%–55%, NM roughly 40%–45%. As a foundry, with massive Capex (USD tens of billions per year on tools and fabs), margins are slightly below NVIDIA, but profitability is formidable after 2nm and CoWoS price hikes.
- **ASML (ASML) (steady rent collector):** GM near 50%–52%, NM about 30%. ASML is heavy industry with long R&D and manufacturing cycles; although each unit sells for USD hundreds of millions, capacity and deliveries are linear and steady.
- **Broadcom (AVGO) (hidden money printer):** its custom ASIC and networking businesses face almost no direct rivals, with GM as high as 65%–70%, NM above 40%, and cash flow extremely healthy.
- **AMD (AMD) (catching up):** its AI accelerators (MI series) carry higher GM, but macro headwinds in traditional PC and general-purpose server markets plus heavy software R&D spend leave overall NM (~15%–20%) the lowest among the Big Five.

3.4 Strategic Dependency/Competition Relationship

If global AI infrastructure faces a crisis, the dependency/competition among the Big Five is as follows:

- **TSMC (TSM US / 2330 TT) and ASML (ASML) (apex layer):** they are the physical ceiling of AI progress. If ASML doesn't ship tools, TSMC can't expand; if TSMC stops, NVIDIA (NVDA), AMD (AMD), and Broadcom (AVGO) run out of supply overnight.

- **NVIDIA (NVDA) (ecosystem ruling layer):** controls the algorithm and application layers. Though it relies on TSMC, it sets the pace and direction of global AI innovation.

- **Broadcom (AVGO) (network hub layer):** bottlenecks AI Scaling. Without Broadcom's high-speed networking, NVIDIA's GPUs—no matter how many—are isolated islands that can't interconnect.

- **AMD (AMD) (strategic balance layer):** the weakest niche, but indispensable—the only weight that checks NVIDIA's pricing power.

3.5 Future growth potential comparison (2H26 outlook)

- **NVIDIA (NVDA):** Massive base; seeking a second S-curve. As the first-phase hardware infrastructure buying spree peaks in 2026–2027, explosive growth in pure hardware sales will decelerate. NVIDIA's future upside hinges on successful commercialization of its NIM (AI software services) and physical AI (sovereign AI, embodied intelligent robots).

- **TSMC (TSM US / 2330 TT):** Exceptionally resilient. As long as global demand for AI compute grows, whether NVIDIA wins, AMD wins, or cloud vendors' in-house chips win, all orders ultimately flow to TSMC. Future growth is driven by the full ramp of 2nm and multi-fold expansion of advanced packaging (CoWoS) capacity.

- **ASML (ASML):** High visibility, but limited upside. Backlog extends years out. Growth over the next few years entirely depends on hard procurement demand from global fabs (TSMC, Intel (INTC), Samsung Electronics (005930 KS)) at 2nm and below for next-gen High-NA EUV (high numerical aperture lithography systems).

- **Broadcom (AVGO):** As hyperscalers (Alphabet (GOOGL), Meta (META), Amazon (AMZN)) aggressively develop in-house AI chips to cut costs, Broadcom's custom ASIC business should inflect over the next 3–5 years; meanwhile, Ethernet's comeback against InfiniBand in 10,000-GPU clusters unlocks major potential for Broadcom's networking chips.

- **AMD (AMD):** Small base, big runway. NVIDIA currently takes the lion's share in AI compute. If AMD can, via software (ROCm ecosystem) optimization, claw 10%–15% share from NVIDIA, its AI revenue and profit can scale by multiples.

4. The Four Mega-Bottlenecks

In this 20-layer global AI supply chain, four layers currently face extremely lethal Mega-Bottlenecks that capital alone cannot easily solve.

Summary of the four core bottlenecks

Bottleneck layer	Bottleneck essence	How to overcome	Estimated time to overcome	Biggest blow-up risk
Layer 1 / Layer 9 (Energy power and power transmission & distribution equipment)	Nuclear plants exist, but power can't reach compute centers (physical equipment shortage)	Bypass the grid with self-built microgrids; source from overseas supply chains	2029 - 2031	Supply-chain disruption caused by trade protectionism
Layer 5 / Layer 7 (Packaging & Memory)	Wafers sufficient, but assembly speed and memory output constrain compute shipments (CoWoS and HBM yield bottlenecks)	TSMC (TSM US / 2330 TT) packaging capacity expanding by multiples; ASML (ASML) tools skewed toward memory	2027	Taiwan Strait geopolitical risk
Layer 17 (Data platform)	High-quality Internet text is exhausted; enterprise internal data is messy	High-fidelity synthetic data; vector knowledge graph plugins	2028 - 2029	Model inbreeding/degeneration; data copyright litigation
Layer 20 (Governance and regulation)	EU AI Act hard-landing in Aug 2026; black-box models hit a compliance wall	Commercialized "mechanistic interpretability" tools; compliance gateways	2028	Compliance shock causes broad delays in enterprise AI deployments

Source: KGI

4.1 Power and Transmission

Nuclear plants exist, but power cannot reach compute centers.

- **Current core issue: Layer 1 (Energy and generation) and Layer 9 (Transmission and distribution equipment).** Substation high-voltage transformers (Transformers) have lead times stuck at 128–160 weeks (~3 years). Grid expansion is severely lagging; nearly half of planned US data centers—and some in Europe with 8–10 year interconnection queues—are delayed or aborted due to no available power.

- **How to overcome? (1) "Behind-the-Meter Gas Microgrids":** Cloud providers skip local grid queues, building distributed natural-gas turbines adjacent to data centers to deliver temporary primary power within 18–24 months. **(2) Modular transformer systems and supply chain reconfiguration:** Western buyers bypass traditional suppliers, aggressively ordering from Chinese transformer majors controlling 60% of global capacity and Japanese/Korean vendors, using modular assembly to cut deployment cycles to 18 months.

- **How long to resolve?** 3–5 years (mitigation expected 2029–2031). Heavy-industry plant capacity expands very slowly, requiring long cycles to digest backlog.

- **Key risk?** Geopolitical trade barriers. If the US and Europe restrict imports of heavy grid equipment from Asian countries for security/geopolitical reasons, their domestic transformer shortage will be extended indefinitely, leaving hundreds of billions of AI assets completely stranded (Stranded Assets).

4.2 Packaging and Memory

Wafers are sufficient, but assembly speed and memory output constrain compute shipments.

- **Current core issue: Layer 5 (advanced packaging) and Layer 7 (HBM memory).** AI chip performance now heavily depends on packaging (CoWoS). TSMC (TSM US / 2330 TT)'s advanced packaging capacity will remain in "extreme allocation (Allocation)" in 2026. Moreover, as demand for HBM3e/HBM4 consumes the entire DRAM industry, HBM wafer Yield Rate remains low, preventing leading-edge GPU shipments due to missing HBM stacks.

- **How to overcome? (1) Packaging capacity spillover and aggressive fab expansions:** TSMC (TSM US / 2330 TT) is outsourcing some non-core packaging steps to leading OSATs (e.g., ASE Technology (3711 TT), Amkor (AMKR)), while its own new advanced packaging fabs will double capacity consecutively over the next two years. **(2) Reallocation of ASML (ASML) equipment:** 2026 Q1 financials show over 51% of ASML's EUV scanners are being directed to Korean memory majors (SK Hynix (000660 KS), Samsung Electronics (005930 KS)), using Low-NA EUV to force-break HBM yield bottlenecks.

- **How long to resolve?** 1.5–2 years (mitigation expected by end-2027). As OSAT capacity comes online and HBM yields improve, the supply gap will gradually narrow.

- **Key risk?** High concentration geopolitical risk in the Taiwan Strait. 99% of top-tier AI advanced packaging is located in Taiwan; any regional conflict would paralyze the global AI supply chain completely within 24 hours.

4.3 Data Platform

High-quality internet text is exhausted; enterprise internal data is a tangled mess.

- **Current core issue: Layer 17 (Enterprise Data Platform).** Public, high-quality human text data has been consumed by frontier models by 2026 (Data Exhaustion). Enterprise AI adoption is slow because historical data in internal legacy systems is extremely dirty, chaotic, and scattered (fragmented ERP, CRM systems); feeding it directly to large models triggers catastrophic business "hallucinations".

- **How to overcome? (1) Physics-level synthetic data pipelines (Synthetic Data):** leverage physics engines and mathematically verifiable systems to generate vertical-industry synthetic data, manufacturing tens of trillions of high-fidelity tokens**. (2) Fuse vector databases with deterministic Knowledge Graphs (RAG):** stop expecting general-purpose large models to understand everything; build robust Knowledge Graphs inside the enterprise as AI "external brains" to translate messy data into structured context.

- **How long to overcome? 2**-3 years (estimated 2028-2029). Data governance and enterprise IT architecture overhauls are a long, painful war of attrition.

- **Key risks?** Data Poisoning and copyright litigation. Improper synthetic data generation can trigger Model Collapse; meanwhile, tightening global jurisdiction over privacy data makes legal landmines easy to step on.

4.4 Governance and Regulation

EU AI Act hard-lands in August 2026; black-box models face a compliance wall

- **Current core issue: Layer 20 (Governance, Security & Regulation**). The EU AI Act's stringent compliance requirements for high-risk AI systems take effect in August 2026. The Act mandates Explainability. Yet deep learning models are non-decomposable mathematical black boxes; if enterprises cannot explain to regulators why a model produced a specific conclusion, they face hefty fines.

- **How to overcome? (1) Commercialize Mechanistic Interpretability tools:** the industry is racing to convert frontier safety science into tools that, like decompiling hacker programs, "translate" activation states across hundreds of billions of parameters inside neural networks into human-readable logical circuits. **(2) Sovereign clouds and Compliance Gateways:** add real-time auditing filters around large models, and at the compliance gateway forcibly cut off or intercept non-compliant or unverifiable logical outputs.

- **How long to overcome?** 2-4 years (a severe compliance shock is expected in 2H26).

- **Key risks?** Innovation stagnation and legal fragmentation. If extremely stringent explainability requirements are not solved technically, Fortune 500 firms, due to compliance risk, will fully pull down or indefinitely delay commercialization of frontier AI applications.

5. The Breakdown of Global USD 7.6 trillion Capex

The global AI industry is entering the largest infrastructure build supercycle in business history. From 2026 to 2031, cumulative AI-related Capex is projected to reach USD 7.6 trillion. In 2026 alone, the combined Capex of the five global cloud hyperscalers has already surpassed USD 600 billion. The essence of today's hundreds of billions in Capex is Big Tech using high-margin software cash to forcibly lock up the next 10 years of global physical infrastructure (power grids, nuclear power, data center land).

Capex split by business allocation

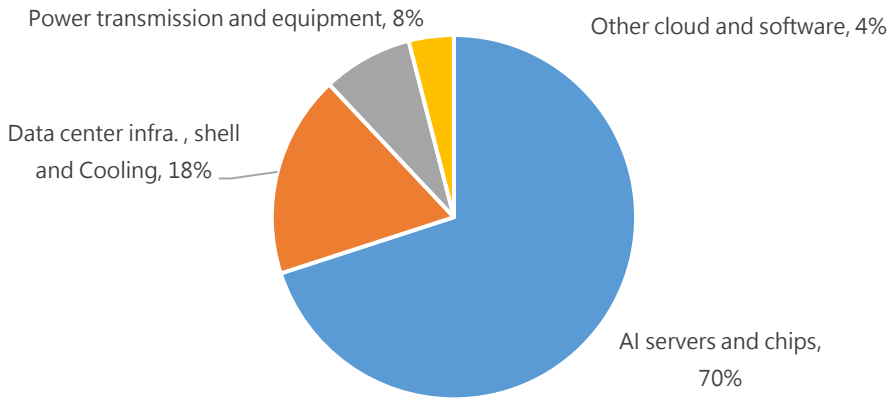

Source: KGI

Capex split by country

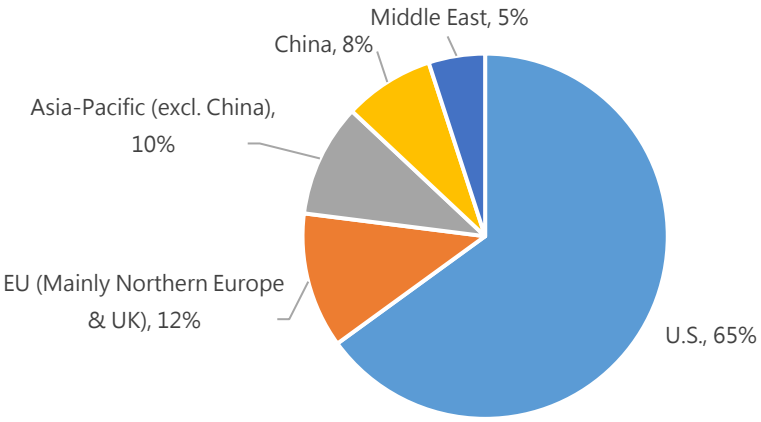

Source: KGI

5.1 Global Capex split by core business allocation

In 2026, the "capital intensity" of Big Tech procurement shifted significantly. Money no longer pours solely into pure chip design; it tilts heavily toward physical infrastructure and power/cooling. The Capex allocation is as follows:

- **70% — AI servers and high-bandwidth silicon (IT Equipment & Silicon):** primarily to purchase NVIDIA (NVDA) Blackwell/Rubin architecture systems, AMD (AMD) MI400 series, and in-house ASICs (e.g., Alphabet (GOOGL) TPU). Within this, high-bandwidth memory (HBM3e/HBM4) and optical networking (1.6T optical modules) have surged, taking nearly 30% of this sub-segment.

- **18% — Data center physical buildings, enclosures, and cooling (Real Estate & Cooling Systems):** to build heavy-duty data centers capable of supporting 100kW+ per rack. Includes civil works and Direct-to-Chip and immersion cooling systems.

- 8**% — Power Delivery & Grid Infrastructure:** Given severe grid constraints, cloud vendors are throwing cash at heavy industrial gear—procuring large high-voltage transformers, switchgear, voltage regulation modules (VRM), and investing in microgrids and energy storage systems.

- **4% — Traditional Cloud & Other:** Spend on non-AI CPU servers and standard office networking hardware has been squeezed to record lows.

5.2 Global Capex distribution by region/country

The geographic layout of AI compute is shifting from “latency-driven (near major cities)” to “power- and policy-driven (go where the power is)”.

- **65% — United States:** still overwhelmingly dominant. The U.S. has the world’s largest tech ecosystem, relatively abundant land, and a flexible power market. CapEx is flowing to the Midwest (e.g., Ohio, Iowa) and Northern Virginia, with numerous “zero-carbon/nuclear” AI towns spanning tens of thousands of acres rising rapidly.

- **12% — Europe (primarily Nordics and the UK):** With core cities like London and Amsterdam facing 8–10-year grid queue bottlenecks, capital is shifting to Iceland, Norway, Sweden—leveraging abundant geothermal/hydro resources and naturally cold climates to host AI liquid-cooling clusters.

- **10% — Asia-Pacific (ex-China mainland):** Mainly into Malaysia (Johor Bahru), Indonesia, and areas around Singapore—emerging as new Southeast Asia AI compute hubs; additionally, Japan and Korea, given strong semiconductor supply chains and sovereign AI demand, are seeing substantial local build-out spend.

- **8% — China mainland**: Led by the three major operators, Alibaba (BABA / 09988 HK), Tencent (00700 HK), Baidu (BIDU / 09888 HK), and Huawei. With leading-edge GPUs blocked by geopolitics, domestic CapEx is

concentrated on ten-thousand-card-scale domestic compute clusters (e.g., Huawei Ascend series) as local alternatives, and on expanding nodes along the "Eastern Data, Western Computing" corridor.

- **5% — Middle East**: Saudi Arabia and the UAE (via sovereign entities like G42) are committing USD tens of billions, leveraging ultra-cheap solar PV and deep sovereign wealth to build Middle Eastern AI compute hubs.

5.3 Construction timeline and payback period

The arrival of the AI heavy-industry era has caused a severe disconnect between project "physical cycles" and "financial cycles":

Project phases	Construction/delivery timeline	Reasons for the long construction/delivery timeline	Investment payback period (Breakeven)
Data center civil works and cooling retrofits	12 — 18 months	The building's civil works are fast, but modern racks are extremely heavy; floor load reinforcement and complex plumbed liquid-cooling commissioning extend timelines.	7 — 9 years (a long-cycle, low-risk real estate/infrastructure ROI profile)
IT hardware asset depreciation cycle	Immediate deployment (delivery requires queuing)	Chips can be powered on to run models as soon as they reach the data hall. However, Blackwell and other chips are in extreme shortage; order-to-delivery lead time is several months.	2 — 3 years (given high compute rental rates, if fully utilized, hardware assets depreciate fast but pay back extremely quickly)
Grid interconnection / transmission equipment build-out	4 — 7 years (up to 10 years in some European cities)	The biggest bottleneck. Lead time for substation high-voltage transformers exceeds 3 years; administrative approval and physical construction for gigawatt-class transmission lines with local grids are extremely lengthy.	Uncontrollable (before power-on, all upfront land and construction spend is "sunk cost")

Source: KGI

5.4 Side-by-side comparison of ROI and risk

1). Which sub-sectors/projects have the highest returns? (The Best ROI)

- **In-house custom AI chips (Custom ASIC):** Alphabet (GOOGL) (TPU) and Amazon (AMZN) (Trainium 2) have the highest ROI in this area. By avoiding the hefty monopoly premium paid to NVIDIA (NVDA), their AI inference services using in-house chips can deliver GM more than 30% higher than directly using NVIDIA GPUs.

- **Bottleneck hardware (Power & Cooling Rollups):** Invest in producers of high-voltage transformers and high-voltage breakers (e.g., Ayr Energy and others that shorten lead times via outsourcing), and in companies directly supplying high-efficiency liquid-cooling pumps. When the whole industry is queuing, these firms enjoy very high premiums and near 40% NM with highly certain returns.

2). Which sub-sectors/projects carry the highest risk? (The Highest Risk)

- **Vacant hyperscale data center real estate (Stranded Assets):** The biggest risk is “shell built, but no power”. In London, Frankfurt, and other regions, many real estate PE funds blindly piled into building data center shells but failed to obtain grid interconnection permits, turning assets worth USD billions into “dead assets” that generate no cash flow over the next 5 years.

- **High-premium generic compute leasing (Tier-2 Neoclouds):** Some second-tier compute cloud companies (Neoclouds) that issued large amounts of high-yield bonds and pledged GPUs to scramble for chips from NVIDIA (NVDA). As the AI market shifts by end-2026 from “blindly training models” to “prioritizing inference cost,” if NVIDIA cuts prices or customers switch to higher-ROI in-house ASIC/Ethernet clusters, these highly levered compute leasing companies will face major bad-debt and bankruptcy/liquidation risks.

6. The Profitability Trends

Global AI ecosystem value chain profit distribution in 2026 is extremely polarized, forming a barbell structure.

Windfall profit nodes	Core representatives	Earnings performance (2026)	Underlying logic
1 Exclusive AI silicon and IP layer	NVIDIA (NVDA), ARM (ARM)	GM: ~74.9%; OM: ~65.6%; NM: ~71.4%	Rigid pricing immunity: enterprises are buying chips—and, more importantly, a compute admission ticket built on the CUDA software ecosystem. R&D costs amortize at scale, driving non-linear profit upside.
2 Monopoly-grade AI gatekeeper memory layer	SK Hynix (000660 KS), Micron (MU)	GM: 53%-58%; OM: 35%- 42%; NM: 28%-34%	Hard necessity to break the "memory wall": HBM4 transitions from a traditional commodity into a highly customized, ASIC-like process. Large buyers prepay 18 months in advance to lock capacity, ending price wars; profit margins approach those of software companies.
3 High-moat B2B SaaS and AI assistants	Microsoft (MSFT), Palantir (PLTR)	GM: 70%- 80%; OM: 40%- 45%; NM: 32%-36%	Exceptionally low marginal cost: Microsoft deeply embeds AI into the Office 365 ecosystem, with underlying compute (Azure) under its own control. The marginal cost of raising enterprise subscription fees is minimal; GM flows cleanly into NM.
Bleeding segments	**Key names**	**Profitability (2026)**	**Underlying logic**
1 Pure-AI independent foundation model layer	OpenAI, Anthropic, AI startups	GM: 25%- 40%; OM: severe losses (<-50%); NM: severe losses	A "sandwich layer" squeezed from both upstream and downstream: upstream, they pay sky-high compute rent to cloud giants (ultimately paying NVIDIA (NVDA)); downstream, they face a global API price-slashing war; revenue cannot cover compute and talent costs.
2 Asset-heavy, high-leverage second-tier compute clouds	CoreWeave; Lambda Labs	GM: 65.5%; OM: -6.9%; NM: -35.6%	Devoured by depreciation and interest: funded by debt secured on GPUs to acquire assets. While the book GM on rented compute is high, once powered on they face massive equipment depreciation, interest, and lease expenses; NM swings deeply negative.
3 Low-barrier AI server assembly and integration	Super Micro Computer (SMCI); Foxconn Industrial Internet (601138 CH)	GM: 6.3%- 9.5%; OM: 3%- 5%; NM: 1.5%- 3%	A "heavy-lifting" business lacking core IP: no proprietary chip IP, only hardware OEM assembly. Due to cutthroat competition, forced to slash margins to grab orders, resulting in revenue exceeding 100 billion and razor-thin net profit.

Note: GM = gross margin; OM = operating margin; NM = net margin. Financial data reflect the latest 2026 dynamics and actual disclosures;

Source: KGI

2026–2030 global AI industry margin trends			
Core trends	**Evolution path and timeline**	**Key catalysts (Catalysts)**	**2028–2030 profitability outlook**
1 Upstream chip/IP layer windfall profits peak and retreat	2026 reaches the windfall-profit ceiling; expected to normalize back to a rational high level after 2028.	1. Foundation models gradually complete the "infrastructure-phase training," shifting compute focus to "inference." 2. Cloud giants' in-house ASIC (e.g., Alphabet (GOOGL) TPU, Amazon (AMZN) Trainium 2) and AMD ramp broadly, breaking the monopoly.	Upstream hardware players such as NVIDIA (NVDA) are expected to see gross margin adjust from the 75% physical ceiling to a rational 60%–65% range, with premium ceded downstream.
2 Compute leasing market undergoes brutal liquidation	2026–2027 see a wave of bankruptcies; after the industry purge, surviving leaders turn profitable.	Highly leveraged, debt-dependent second-tier compute clouds will suffer cash-flow breaks as equipment depreciation costs peak, ushering in a "darkness before dawn" and capacity reshuffle.	Surviving top compute clouds, after equipment depreciation runs off (post-2028), will see GM recover from the 2026 trough of 8%–15% to 15%–25%, with NM turning profitable.
3 Ultimate profit harvesters shift fully to the application layer	2028–2030 ushers in a golden harvest, with the profit center shifting from "sellers of shovels" to "power users".	Underlying compute costs are collapsing by 90% as technology iterates, sharply lowering software vendors' operating costs. Vendors that control end-to-end enterprise workflows and exhibit strong B2B stickiness command exceptional premium pricing power.	Vertical AI and SaaS vendors will enjoy ultra-low compute costs and high subscription premiums; net margin could see non-linear takeoff, sprinting to 35%+.

Source: KGI

The analysis below leverages Gross Margin (GM), Operating Margin (OM), and Net Margin (NM) to deeply deconstruct the 2026 global AI ecosystem's "Profit Powerhouses" and "Bleeding Centers," and to project future margin trends.

6.1 Profit Powerhouses

1). Exclusive AI silicon and IP layer

- Key representatives: NVIDIA (NVDA), ARM (ARM).
- Financial performance (2026): GM: 74.9% | OM: 65.6% | NM: 71.4% (NVIDIA 2026 Q1 / FY27 Q1 actual data; NM was elevated by other non-operating income, but OM reached 65%).
- Underlying drivers: NVIDIA's Blackwell Ultra and next-generation Rubin architecture systems enjoy industry-wide "pricing immunity." Enterprises are buying not just chips, but a "compute admission ticket" anchored in the CUDA software ecosystem. The formidable software moat means once R&D is scaled, OM and NM exhibit explosive non-linear expansion.

2). Monopoly-grade AI gatekeeper memory (Memory Layer)

- Key representatives: SK Hynix (000660 KS), Micron (MU).
- Financial performance (2026): GM: 53% — 58% | OM: 35% — 42% | NM: 28% — 34%.
- Underlying drivers: The "memory wall" has erupted; HBM4 has transformed from traditional memory into a highly customized ASIC process, with major buyers securing capacity via 18-month prepayments.

This allows memory giants to leave behind the old “price-cut grinder” fate, pushing margins toward first-tier software companies.

3). High-moat B2B SaaS and AI assistants (Ecosystem Application)

- Key representatives: Microsoft (MSFT) Copilot, Palantir (PLTR).
- Financial performance (2026): GM: 70% — 80% | OM: 40% — 45% | NM: 32% — 36%.
- Underlying drivers: Microsoft has deeply embedded AI into Office 365 used by over a billion white-collar workers globally. Because the underlying compute (Azure) is in-house, the marginal cost of charging enterprises an extra USD 30 per month for AI subscriptions is minimal, enabling near-perfect translation from gross margin to net margin.

6.2 The Bleeding Centers

1). Pure AI standalone large model layer (Foundation Model Frontier)

- Key representatives: OpenAI, Anthropic, and various open-source/closed-source startups.
- Financial performance (2026E): GM: 25%—40% | OM: negative (even below -50%) | NM: severe losses.
- Underlying reason: they sit in the most brutal “squeezed middle.” Upstream, they must pay extremely expensive rent to AWS/Microsoft (MSFT) (indirectly footing NVIDIA (NVDA)’s high GM) to keep trillion-parameter model training and inference running; downstream, a 2026 API price war (Alibaba (BABA / 09988 HK), Alphabet (GOOGL), etc. slashing Token prices) makes model revenue unable to cover their high compute and top-talent costs.

2). Heavy-asset, high-leverage Tier-2 compute leasing clouds (Tier-2 Neoclouds)

- Key players: CoreWeave, Lambda Labs, and other compute clouds highly reliant on debt financing.
- Financial performance (2026): GM: 65.5% (optically high) | OM: -6.9% | NM: -35.6% (CoreWeave 2026 Q1 disclosure shows deep net losses).
- Underlying reason: these firms pledged GPUs to borrow USD tens of billions to lease data centers and stack assets. While compute rental GM looks great (65%), once they power on, heavy depreciation, interest expense, and long-term data center Lease Costs instantly consume cash, sending OM and NM plunging into the red.

3). Low-barrier AI server assembly and system integration (Server Integration)

- Key players: Super Micro Computer (SMCI), Foxconn Industrial Internet (601138 CH).

- Financial performance (2026): GM: 6.3%—9.5% | OM: 3%—5% | NM: 1.5%—3% (SMCI's GM in 2026 was severely compressed to below 10% in H1 this year).

- Underlying reason: they do the AI era's heavy lifting. Although Blackwell orders pack server-makers' schedules, they lack core silicon IP—just bolting NVIDIA (NVDA) chips, SK Hynix (000660 KS) memory, and liquid-cooling plumbing into racks. With fierce competition, they must squeeze GM to win orders—classic huge throughput, razor-thin profit.

6.3 Future margin trajectory (2026–2030)

Trend 1: Supernormal profits at the upstream chip/IP hardware layer will peak and roll off in the next 2–3 years.

- 2026 marks the physical ceiling for NVIDIA (NVDA) achieving 75% GM.

- As LLM "infrastructure training" largely completes in coming years and global compute tilts to inference, cloud giants' in-house ASICs—Alphabet (GOOGL) TPU, Amazon (AMZN) Trainium 2—and AMD's ramp will drive NVIDIA's GM to revert after 2028 to a rational 60%–65%, with premium ceded to downstream.

Trend 2: Compute leasing clouds will see margin recovery only after a brutal clearing.

- Debt-fueled Tier-2 compute clouds face the "darkness before dawn." With depreciation costs peaking this year, 2026–2027 should see a wave of highly leveraged compute lessors go bankrupt or get acquired as cash flow breaks.

- After the shakeout, leading survivors, once assets fully depreciate in 2028, should see GM recover from the 8%–15% trough in 2026 to 15%–25% in 2028–2030; NM turns positive.

Trend 3: The ultimate profit "harvester" will shift fully to the application layer (SaaS).

- By 2028–2030, as underlying compute costs collapse 90% via iterative tech progress, the true golden era of application-layer explosion arrives. Vendors owning end-enterprise workflows and sticky B-side relationships will enjoy ultra-low compute costs and outsized AI subscription premium. The AI profit center shifts epicly from "pick-and-shovel sellers (NVIDIA (NVDA)/memory majors)" to "power users (SaaS/vertical AI)."

- NM in 2028–2030 has room to surge, sprinting to 35%+ in a golden harvest phase.

7. AI Model Comparison: US/CN

In the 2026 global AI landscape, US-China competition among AI models has entered the "deep-water" stage. The earlier absolute US technical lead has faded; China shows exceptional cost advantages in parts of the stack. According to Stanford University's latest AI Index Report 2026 (Stanford AI Index Report 2026), the performance gap between top US and China AI models has fallen from double digits to about 2.7%.

- **US:** Leveraging the world's largest USD capital, the top-tier chip supply chain, and hyperscale compute where brute force delivers miracles, it remains aggressive at the AGI frontier and in multimodal ecosystems.

- **China:** Despite chip constraints and no absolute capital advantage, through extreme algorithmic squeezing, ruthless engineering optimization, and a pragmatic free/open-source strategy, it has driven AI inference cost to near-zero, showing strong aggression in productivity democratization and industrial deployment.

7.1 Model Performance and Technical Route

Extreme compute vs extreme algorithms

Dimension	US model ecosystem	China model ecosystem
Top representatives	OpenAI (GPT-5/o3)、Anthropic (Claude 4.6)、Google (Gemini 2.5)	DeepSeek (V4/R1), Jieyue Xingchen (Dola-Seed 2.0), Moonshot AI (Kimi K2.6)
Technical route	Heavy-industry route (Brute Force): betting on "brute force delivers miracles," pushing compute and parameter ceilings by building hundreds-of-billions-scale data centers (e.g., Microsoft (MSFT) and OpenAI's Stargate program).	Light-industry route (Extreme Efficiency): focused on "max energy efficiency." Using MoE, multi-token prediction, and FP8 mixed-precision training to squeeze every unit of compute under severe chip constraints.
Core advantages	Full-spectrum multimodality and complex Agents: maintains high barriers in native voice, Sora-grade video generation, and automated cross-software complex workflows.	Deep reasoning, math/coding, and ultra-long context: in math contests (e.g., AIME) and low-level code development, China models achieve parity or even surpass with minimal resource consumption.

Source: KGI

7.2 Business Model and Cost Efficiency

High-ticket premium vs price disruptors

- **US: premium closed-loop.** Because US vendors (e.g., Anthropic, OpenAI) must cover heavy R&D expenses, stringent Red-Teaming security review costs, and high compute depreciation, flagship models (e.g., GPT-5.4 or Claude Opus) typically price APIs at USD 2.5–15/mn Token. They monetize via high-ticket enterprise-grade solutions and consumer subscriptions (e.g., ChatGPT Pro).

- **China: the ultimate price destroyer (Memflation reversal).** Represented by DeepSeek, China's foundation models (V4 or R1 inference) are priced at USD 0.28-0.55/mn Token for APIs in 2026, 20-30x cheaper than comparable US models on output and inference. This floor price has galvanized deployments by independent developers and

SMEs worldwide, ushering foundation models into an era of compute democratization.

7.3 Innovation Ecosystem

Closed-source gatekeepers vs open-source ecosystem

- **US: closed-source dominance and oligopoly.** The top-tier AI capabilities (GPT series, Claude series) are firmly controlled by closed-source vendors. These vendors are deeply tied to Hyperscalers such as Microsoft (MSFT), Amazon (AMZN), and Alphabet (GOOGL), forming exclusive ecosystem camps. Open-source forces (e.g., Meta (META)'s Llama 4) are strong but typically trail the very best closed-source models by 4-6 months in absolute performance.

- **China: the comprehensive victory of Open-Weights.** China's foundation models are highly open-source; whether DeepSeek or Alibaba (BABA / 09988 HK) Qwen, the weights of top inference models are largely open to the world. This makes Chinese models the de facto standard for the global open-source ecosystem by 2026, as countless companies pursue local-language deployments and Fine-tuning of Chinese open-weight models, bypassing costly US closed-source APIs.

7.4 Resource Endowments and Infrastructure

Capital/Chips vs Power/Patents

- **US: private capital and high-end chips.** US AI private investment totals USD 285.9 billion in 2025—23x China. Despite strict chip export curbs, US data centers still host the world's densest top-tier GPU compute clusters.

- **China: power infrastructure, applied base, and patent surge.** US infrastructure (e.g., the grid) faces serious expansion bottlenecks; by contrast, China's power reserves are abundant (reserve margin never below 80%), supplying ample, low-cost green energy for AI compute. China accounts for 69.7% of global AI patent filings (~4x the US) and installs 9x as many industrial robots as the US, implying far richer physical deployment soil for Chinese models across industrial, manufacturing, and hardware scenarios.

7.5 Talent and Regulation

Global hub vs manufacturing base

- **Talent flows:** The US remains the ultimate destination for the world's top AI scientists. But dynamics are shifting: due to geopolitics, net inflows of overseas AI talent into the US have fallen 89% since 2017, while China's domestically trained algorithm engineering corps (engineers highly experienced in optimization and distillation) is scaling up, forming a formidable technical base.

- **Regulation and compliance:** US reviews on AI safety, bias mitigation (Responsible AI), and copyright compliance are onerous, slowing iteration to a degree. China's policy focus is more pragmatic—driving foundation-model deployment in the real economy and industrial manufacturing to reduce enterprise production costs.

8. AI Trade Peaked? How to Position?

7.1 Why We Think the AI Trade Has Not Peaked Out?

1) **Today's valuation of 37x vs 45–100x during bubble periods.** In past bubbles, P/E in the affected segments at least surged to 45–72x on a trailing 12-month basis. In Asia in particular, Taiwan and Japan reached peak trailing 12-month P/E of 72x and 100x in 1989. By contrast, the MSCI ACWI Semiconductors sector now trades at 37x trailing P/E.

2) **Fundamentals remain robust with earnings estimates keep moving higher.** Over the past 3 months, corporate EPS revisions have been raised 10% more than the broad market. 12-month forward EPS growth is estimated at 40%, versus 12% for the MSCI AC World Index.

3) **Upstream capacity constraints will take years to resolve.** Supply remains tight because: (A) Long fab build cycles (it takes 3 years to build a wafer fab); (B) TSMC (TSM US / 2330 TT) maintains disciplined Capex management (Capex is only up ~80% vs the pre-GenAI period); (C) Memory capacity can only increase ~30% across 2026–2027; and (D) Capex growth at TSMC and memory makers lags hyperscalers materially

Relative to the TMT bubble, overinvestment is limited: IP equipment and software spending as a share of GDP has risen 17% since 2022, versus a 42% increase during 1994–2000.

4) **TAM expectations have significant room for upward revision.** Assuming semiconductor hardware and software account for 2% of global GDP in 2030, we estimate industry revenue (TAM) at USD 3 trillion. If the 2030 net margin stays at the 2025 level of 35%, the implied P/E is 16x; if the 2030 net margin is steady at 25% (vs. 25% in 2025), the implied P/E is 22x. From this perspective, neither TAM assumptions nor valuations look overly aggressive. As a reference, oil's share of global GDP has averaged 3% historically (spent long stretches between 4% and 5%, peaking at 10%), and AI may warrant a higher share.

5) **Sheer market cap has already capped how far investors can Overweight.** Many funds stipulate that a single-stock position cannot exceed 10% of total portfolio market value. This means these investors now struggle even to reach a Market-weight allocation versus the index. TSMC (TSM US / 2330 TT), Samsung Electronics (005930 KS) and NVIDIA (NVDA) are already facing this issue. The problem is even more acute in Growth mandates.

6) **Future deployable capital equals 60% of the current sector market cap.** About USD 2.3 trillion is flowing into the US equity market via dividends, share buybacks, and cash-funded M&A. Together with up to USD 7.8 trillion still parked in money market funds, the total pool exceeds USD 10 trillion—equivalent to 60% of the sector's current market cap of USD 16.6 trillion.

Market-cap concentration limits how much investors can Overweight

Stocks	Index	% Market Cap
Broadcom	MSCI US	3.5%
ASML	MSCI Europe	4.9%
NVIDIA	MSCI US	8.0%
SK Hynix	MSCI EM	8.7%
ASML	MSCI Europe growth	9.6%
NVIDIA	MSCI ACWI growth	10.0%
SEC	MSCI EM	10.8%
TSMC	MSCI EM	15.6%
Source: UBS, Refinitiv DataStream, KGI		

7.2 How to Position? KGI Global AI 30 Basket

Listed companies breakdown: US: 19; China: 5; Taiwan: 2; South Korea: 2; Netherlands: 1; Japan: 1

Name and Ticker	Investment Thesis	Market Cap (USD mil)	Sales YoY (%) 2026	Sales YoY (%) 2027	EPS YoY (%) 2026	EPS YoY (%) 2027	EBIT Margin 2026	EBIT Margin 2027	ROE 2026	ROE 2027	FCF Yield	EV/Sales	P/E 2026	P/E 2027
(1) Global AI Infra Monopoly														
Nvidia (NVDA.US)	Monopolistic GPU/CUDA; compute dominance; high margins, reasonable valuation.	4,736,908	65.5	81.5	**59.5**	**88.7**	63.4	66.2	92.9	92.6	**2.5**	18.3	41.8	**21.8**
TSMC (TSM.US)	Monopolizes top-tier global wafer fabrication and packaging; high margins, reasonable valuation.	2,256,064	36.5	27.8	**49.2**	**26.7**	57.3	56.9	38.8	37.1	**1.9**	14.3	28.1	**22.1**
Broadcom (AVGO.US)	Monopolizes high-speed networking chips; asset-light, high gross margin, strong free cash flow.	1,802,695	65.6	63.8	**69.4**	**66.2**	65.5	65.0	54.6	59.6	**1.8**	24.5	32.8	**19.7**
Micron (MU.US)	HBM3e/HBM4 breakthroughs; top share contender; geopolitical tailwinds.	1,370,586	238.2	89.8	**764.9**	**107.1**	74.6	81.6	74.5	71.8	**1.9**	14.9	16.9	**8.2**
Samsung Electronics (005930.KS)	Downstream vertical integration across the full value chain; early HBM misjudgment; catching up.	1,251,908	106.6	27.0	**495.6**	**38.5**	51.5	57.1	49.8	44.1	**2.9**	4.7	7.5	**5.4**
SK Hynix (000660.KS)	Global HBM share No.1 (58–61%); holds the vast majority of NVIDIA (NVDA) core orders.	1,241,648	251.9	47.4	**394.9**	**44.8**	76.5	77.8	92.8	62.7	**2.2**	14.3	8.8	**6.1**
ASML (ASML.NA)	Monopolizes EUV lithography tools; relatively oversold; earnings revisions	702,852	19.6	23.2	**28.2**	**33.5**	36.6	40.1	55.5	58.9	**1.4**	18.2	50.3	**37.7**
	Simple Average		**112.0**	**51.5**	**265.9**	**57.9**	**60.8**	**63.5**	**65.6**	**61.0**	**2.1**	**15.6**	**26.6**	**17.3**
	Market Cap Weighted Average		**97.1**	**59.5**	**201.8**	**65.4**	**62.5**	**64.9**	**70.7**	**67.6**	**2.2**	**16.4**	**29.9**	**18.0**
(2) Mega 7														
Nvidia (NVDA.US)	Monopolistic GPU/CUDA; compute dominance; high margins, reasonable valuation.	4,736,908	65.5	81.5	**59.5**	**88.7**	63.4	66.2	92.9	92.6	**2.5**	18.3	41.8	**21.8**
Alphabet (GOOGL.US)	GCP cloud; Gemini best-in-class; search engine and ads monetized	4,175,544	23.2	21.2	**11.1**	**7.4**	40.3	40.6	49.5	38.0	**1.5**	9.8	24.1	**22.4**
Apple (AAPL.US)	AI applications at scale; near-term hardware innovation slowing; resilient free cash flow.	4,041,226	14.9	9.0	**17.5**	**10.2**	32.4	31.9	132.5	108.0	**3.2**	8.8	31.4	**28.5**
Microsoft (MSFT.US)	Azure cloud, OpenAI; generative AI embedded into enterprise software; compounding growth.	2,620,975	16.9	16.7	**25.5**	**13.9**	46.6	46.7	31.1	27.6	**2.8**	8.4	20.6	**18.1**
Amazon (AMZN.US)	AWS cloud growing strongly; e-commerce margin expansion; balanced offense and defense.	2,441,971	15.1	13.2	**14.9**	**12.5**	12.6	14.1	19.8	17.8	**-0.1**	3.4	22.0	**19.6**
Tesla (TSLA.US)	Full self-driving FSD, energy storage and humanoid robots; highest beta.	1,408,847	8.2	15.2	**12.8**	**30.0**	5.3	6.7	6.4	7.6	**0.5**	14.1	200.3	**154.1**
Meta Platforms	AI boosts FB/IG ad targeting and monetization; new-tech acquisition	1,378,034	25.9	19.2	**27.9**	**-4.1**	35.7	34.9	32.1	25.5	**3.5**	6.4	13.5	**14.1**
	Simple Average		**24.2**	**25.2**	**24.2**	**22.7**	**33.7**	**34.4**	**52.0**	**45.3**	**2.0**	**9.9**	**50.5**	**39.8**
	Market Cap Weighted Average		**28.6**	**30.5**	**26.9**	**28.6**	**38.9**	**39.7**	**65.6**	**57.5**	**2.1**	**10.7**	**40.1**	**30.9**
(3) Segment Leaders														
AMD (AMD.US)	Direct competitor to NVIDIA (NVDA); MI300/400 series offers superior price/performance.	868,409	43.7	53.8	**73.8**	**81.0**	26.8	31.7	14.2	21.5	**1.0**	23.0	73.5	**40.6**
Renesas (6723 JP)	Automotive chip giant; AI edge computing; smart-vehicle cycle recovery.	59,093	16.9	9.8	**37.3**	**5.6**	24.7	23.0	14.6	12.2	**3.7**	7.5	20.4	**19.3**
Sandisk (SNDK.US)	Pure-play NAND flash and enterprise SSD exposure; storage supercycle.	345,790	169.0	130.5	**2,095.5**	**191.3**	59.9	78.4	60.5	71.0	**1.3**	26.0	35.6	**12.2**
Marvell (MRVL.US)	Absolute leader in Optical/DSP; core partner on ASICs; high-margin, asset-	246,046	42.1	40.6	**80.9**	**42.4**	35.3	36.3	2.8	12.0	**0.7**	28.4	99.1	**69.6**
Coherent (COHR.US)	Global optical-communication materials giant; asset-heavy with high depreciation; profits surge once mass production ramps.	79,674	21.5	36.2	**54.5**	**52.6**	20.4	22.4	9.9	12.9	**-0.7**	12.3	74.7	**48.9**
Lumentum (LITE.US)	Monopolizes global EMI optical chips; order visibility directly tied to total global compute expansion.	67,061	82.2	85.7	**295.9**	**121.6**	29.1	39.0	32.7	44.8	**0.2**	27.0	105.7	**47.7**
GE Vernova (GEV.US)	Monopolizes gas turbines and grid equipment markets; high-certainty energy exposure with strong orders.	291,687	19.4	14.3	**142.0**	**16.0**	10.2	14.7	39.8	36.1	**2.6**	7.3	52.7	**45.4**
Eaton (ETN.US)	Global giant in electrical management and grid infrastructure; high margins, asset-light, high certainty.	163,036	16.3	10.5	**10.1**	**16.8**	19.8	21.0	23.5	24.5	**2.3**	6.5	31.6	**27.0**
Vertiv (VRT.US)	Absolute leader in liquid-cooling thermal management; dominates the cooling market for NVIDIA (NVDA) Blackwell and other facilities.	125,054	35.6	28.2	**54.0**	**34.0**	23.2	24.5	48.5	46.2	**1.8**	11.6	50.3	**37.6**
Asia Vital Components (3017.TT)	Monopolizes 70% share of AI fans and liquid-cooling key components; tightly tied to NVIDIA (NVDA).	28,337	57.8	35.5	**92.3**	**44.3**	25.0	26.3	58.1	57.1	**4.2**	5.3	24.3	**16.9**
	Simple Average		**50.4**	**44.5**	**293.6**	**60.6**	**27.4**	**31.7**	**30.4**	**33.8**	**1.7**	**15.5**	**56.8**	**36.5**
	Market Cap Weighted Average		**56.9**	**53.4**	**390.2**	**75.9**	**29.7**	**35.6**	**26.8**	**32.0**	**1.4**	**19.3**	**62.6**	**38.5**
(4) Strategic Reshaping														
Intel (INTC.US)	Benefits from US political subsidies for domestic advanced foundry; earnings rebound through the transition.	667,805	11.0	11.5	**157.6**	**43.7**	12.2	15.4	3.6	6.5	**-0.5**	12.7	122.8	**85.4**
Dell (DELL.US)	Co-selling AI servers with NVIDIA (NVDA) ($50bn AI orders in-hand Backlog).	264,563	18.8	51.5	**26.5**	**74.4**	8.7	9.3	-353.7	-422.9	**3.6**	2.1	41.1	**22.8**
	Simple Average		**14.9**	**31.5**	**92.1**	**59.1**	**10.4**	**12.3**	**-175.1**	**-208.2**	**1.6**	**7.4**	**82.0**	**54.1**
	Market Cap Weighted Average		**13.2**	**22.8**	**120.4**	**52.4**	**11.2**	**13.6**	**-97.8**	**-115.4**	**0.7**	**9.7**	**99.6**	**67.7**
(5) China AI														
Tencent (700.HK)	WeChat ecosystem + games = stable cash flow; AI ads monetized with precision; Tencent (00700 HK) Cloud re-rating.	478,767	10.3	9.4	**8.3**	**11.0**	32.6	33.2	19.3	18.6	**6.2**	4.3	11.9	**10.7**
Alibaba (BABA.US/9988.HK)	Alibaba (BABA) Cloud/Qwen benefits from China's AI compute boom; legacy e-commerce as defense.	228,240	2.7	9.6	**-59.0**	**59.4**	6.3	8.1	8.0	8.7	**-3.2**	1.1	20.1	**15.1**
Knowledge Atlas / Zhipu AI	China LLM "new seven little dragons"; flagship domestic generative AI; potential inclusion in the Hang Seng TECH Index.	125,080	346.2	173.0	**-32.8**	**6.2**	-147.0	-55.6	46.3	554.2	**N/A**	1,172.7	N/A	**N/A**
Victory Giant (2476.HK/300476.SZ)	Core leading supplier of high-end PCB/HDI boards for NVIDIA (NVDA) GPUs and AI servers.	46,430	75.1	70.6	**95.8**	**77.6**	31.0	32.3	30.8	38.1	**-1.0**	15.8	33.2	**18.7**
Baidu (BIDU.US/9888.HK)	Cloud revenue +79%; in-house Kunlunxin chip spun off for listing; Apollo Go robotaxi autonomous driving.	35,383	2.7	7.7	**-5.6**	**18.1**	10.2	11.5	5.3	5.9	**-4.1**	1.0	14.0	**11.9**
	Simple Average		**87.4**	**54.1**	**1.3**	**34.4**	**-13.4**	**5.9**	**21.9**	**125.1**	**-0.5**	**239.0**	**19.8**	**14.1**
	Market Cap Weighted Average		**57.4**	**34.9**	**-10.3**	**26.1**	**0.5**	**13.9**	**20.2**	**89.9**	**2.2**	**163.9**	**13.5**	**10.8**
	KGI AI 30 - Simple Average		**62.0**	**41.1**	**169.9**	**45.7**	**27.6**	**33.0**	**26.5**	**39.9**	**1.5**	**50.8**	**45.1**	**31.3**
	KGI AI 30 - Market Cap Weighted Average		**53.0**	**36.3**	**118.2**	**38.6**	**42.4**	**44.3**	**55.4**	**50.9**	**2.0**	**16.6**	**38.2**	**28.1**

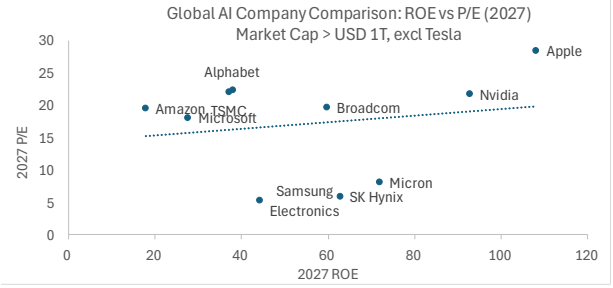

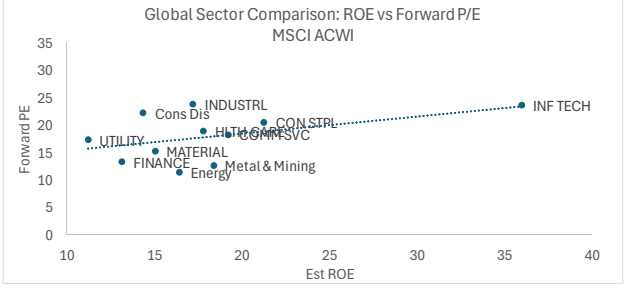

Source: Bloomberg, KGI (As of 26/06/2026)

(1) Global AI Infra Monopoly

- **NVIDIA (NVDA):** Monopolizes GPU hardware and global CUDA software development; compute dominance; the majority of industry Capex will translate into NVIDIA's revenue/profit; high margins, reasonable valuation.

- **TSMC (TSM US / 2330 TT):** Monopolizes foundry production of all top-tier AI compute chips globally; possesses world-leading leading-edge wafer fabrication and CoWoS advanced packaging technology; high margins, low valuation.

- **Broadcom (AVGO):** Dominates high-speed networking silicon for high-end AI clusters; exclusive development partner for custom AI chips (ASIC) for Alphabet (GOOGL) and other tech giants; asset-light, high-GM model sustains robust FCF and high dividends.

- **SK Hynix (000660 KS):** No.1 global market share in high-end AI memory (HBM); NVIDIA (NVDA)'s primary memory partner; multi-year order visibility locked in; highest certainty within the AI compute hardware ecosystem outside of NVIDIA (NVDA).

- **Micron (MU):** Leveraging breakthroughs in power-efficient HBM (HBM3e/HBM4), it has entered the top-tier AI compute supply chain, benefiting from a memory supercycle with rising volumes and pricing driven by AI large models' inelastic need for massive, high-bandwidth memory. The only pure-play memory stock in US equities.

- **Samsung Electronics (005930 KS):** Having put in a major bottom after early HBM lag, now at a key inflection point as its next-gen high-end HBM4 memory has cleared qualification with top-tier customers.

- **ASML (ASML):** Monopoly supplier of high-end EUV lithography systems required to manufacture advanced chips—a bottleneck TSMC (TSM US / 2330 TT), Intel (INTC), and Samsung Electronics (005930 KS) cannot bypass when expanding leading-edge capacity; relatively oversold, with earnings estimate revisions lagging.

(2) Mega 7

- **NVIDIA (NVDA):** Monopolizes GPU hardware and the global CUDA software ecosystem; compute dominance; the majority of industry Capex will translate into NVIDIA's revenue/profit; high margins; reasonable valuation.

- **Alphabet (GOOGL):** Has demonstrated top-tier capabilities in GenAI through Gemini, effectively monetizing its search engine and massive ad traffic; valuation is relatively inexpensive within the Magnificent Seven.

- **Apple (AAPL):** Owns the world's stickiest hardware-software ecosystem, the natural gateway for AI applications reaching consumers; even if near-term hardware innovation slows, its robust free cash flow and buyback firepower effectively support share-price performance.

- **Microsoft (MSFT):** Leveraging its deep investment in OpenAI and Azure cloud's dominant edge, it led in seamlessly embedding GenAI into Office and other enterprise software, underpinning strong B2B pricing power and compounding growth.

- **Amazon (AMZN):** Its core profit engine, AWS cloud computing, has seen a strong rebound in growth, while the e-commerce business has delivered significant margin expansion after completing supply-side optimization—making it a balanced recovery play.

- **Meta Platforms (META):** AI is significantly enhancing ad targeting precision and monetization on its FB/IG social platforms, while acquisitions of new technologies (e.g., Manus AI) continue to broaden its hardware and software boundaries.

- **Tesla (TSLA):** Even as pure EVs face cyclical headwinds, its core value proposition has fully pivoted to Full Self-Driving (FSD), energy storage, and humanoid robots—making it the highest-beta future-tech stock among the Magnificent Seven.

(3) Segment Leaders

- **AMD (AMD):** NVIDIA (NVDA)'s direct competitor, staying close with the MI300/400 series; the bet is on capturing spillover compute orders from a nonlinear inflection and achieving higher share-price beta.

- **Renesas (6723 JP):** Holds formidable moats in the automotive supply chain; amid a JPY asset re-rating, a slow-bull play with dual exposure to Edge AI and intelligent-vehicle cycles, offering solid downside protection.

- **SanDisk/Western Digital (SanDisk - SNDK US):** A classic "cycle upturn + AI must-have" high-beta stock; explosive data center demand for high-end, high-capacity SSDs is propelling it into an earnings boom phase.

- **Marvell (MRVL):** In the high-speed optical module/DSP market (within AI servers, electrical signals are converted to optical for transmission and managed via DSP chips for control and error correction), Marvell and Broadcom (AVGO) form a duopoly. It is also a core partner for Big Tech's in-house ASIC AI chips. High GM, asset-light.

- **Coherent (COHR):** One of the world's largest suppliers of optical modules and optoelectronic components, integrating optical chips,

housings, and interfaces into complete optical modules (Transceivers) sold to data centers. It is also a global giant in wafer materials such as SiC and Indium Phosphide (InP), controlling the most upstream materials processes for optoelectronic conversion. Asset-intensive with heavy depreciation; once 1.6T optical modules reach mass production, net profit should surge nonlinearly.

- **Lumentum (LITE):** Monopolizes global EMI physical photonics chips—the core, highest-barrier light-emitting component in optical modules. Without high-end optical chips, optical modules cannot emit optical signals. As line rates rise from 800G to 1.6T (a 2x step), demand for high-end EML chips and CW (continuous-wave) high-power lasers increases inelastically, roughly doubling. Its de facto monopoly ties order visibility directly to the global compute capacity buildout.

- **GE Vernova (GEV US):** An independently listed pure-play power heavy equipment name spun off from GE; owns the world's most advanced, highest-efficiency HA-class gas turbines (sub-one-minute fast start). It helps Hyperscalers avoid waiting 3–5 years for traditional grid expansions by directly procuring utility-scale generator sets to bring data centers online faster; a high-visibility energy name (backlog extends several years out).

- **Eaton (ETN):** Focused on medium- and low-voltage power distribution and in–data center electrical architecture (transformers stepping high voltage down to equipment-ready power, UPS, and intelligent rack PDU). A prime play on the AI energy bottleneck: high-margin, asset-light (no heavy generator-set manufacturing; focused on last-mile distribution hardware near servers and power management software); high visibility (US data center expansions have upgrade orders for distribution systems and transformers booked out for years).

- **Vertiv (VRT):** The absolute leader in high-temperature liquid-cooling thermal management, benefiting as high-power AI chips from NVIDIA (NVDA) and others trigger massive, inelastic demand for data-center liquid cooling; high fundamental visibility (independent of competition among LLMs—so long as compute capacity expands, customers must order the company's products).

- **AVC (3017 TT):** Leveraging a dual-track strategy of 3D VC and cold plates to become the thermal leader among Taiwan-listed stocks; tightly aligned with NVIDIA (NVDA)'s latest server architecture upgrade cadence, enabling faster technology deployment than US-listed peers.

(4) Strategic Reshaping

- **Intel (INTC):** In a major transformation, benefiting from policy-driven US subsidies for domestic advanced foundry and from a

commercial “AI PC” refresh cycle that offers strong margin recovery leverage; earnings are poised to rebound.

- **Dell Technologies (DELL):** The go-to hardware partner for traditional government and enterprise customers to operationalize AI, leveraging a deep partnership with NVIDIA (NVDA) to sell AI servers. In a high-valuation tech environment, its robust, stable cash flows, planned share buybacks, and a 27% payout ratio, together with a reasonable valuation range, offer investors attractive risk-reward.

(5) China AI stocks

- **Tencent (00700 HK):** Leveraging WeChat’s uniquely sticky ecosystem and massive cash flows, accelerating two-sided monetization across B/C via overseas gaming expansion and the Hunyuan LLM; a value stock with strong anti-cyclical resilience and AI-driven incremental upside, trading at an attractive valuation.

- **Alibaba (BABA / 09988 HK):** Core e-commerce bottoming out with profits and market share recovering; rapid overseas cross-border e-commerce expansion; Alibaba Cloud set to fully benefit from the surge in China enterprise AI compute; low valuation and ongoing share buybacks provide a high margin of safety; valuation remains cheap.

- **Baidu (BIDU / 09888 HK):** China’s GenAI commercialization leader, reshaping search advertising and cloud via LLMs; autonomous driving has reached scaled commercial operations across multiple cities nationwide; a high-moat tech stock with strong valuation optionality, oversold and cheap.

- **Victory Giant Tech (02476 HK / 300476 SZ):** As the leading PCB supplier to NVIDIA (NVDA) for core AI servers and GPUs, it is a direct beneficiary of the global AI compute arms race; surging orders for high-layer-count, high-density PCBs are driving results; dual listings in Hong Kong and Mainland China support valuation re-rating potential.

- **Zhipu AI (2513 HK):** The Tsinghua-lineage player with the deepest technical roots among China’s “New Seven Little Dragons” of LLMs; GLM full-stack, self-developed foundation model; dual backing from the national team and Big Tech; B2B commercialization across all scenarios; exceptional ecosystem penetration across academia, government/enterprise, and consumer. A China GenAI bellwether with the highest technical certainty and the most robust closed-loop commercialization.

7.3 Industry Upside Catalysts

If the AI value chain can smoothly relay from "hardware heavy industry" to a "software supercycle," shares still have room to run. Investors need to see the following:

1) Agentic AI delivering B2B cost-down and efficiency gains with hard "financial proof"

- 2026 is the rollout year for Agentic AI. The market will no longer be satisfied with how many people registered for ChatGPT; it will look to whether Microsoft (MSFT) Copilot, Palantir (PLTR), or Salesforce (CRM) AI agents can truly cut Fortune 500 operating costs by 20% or deliver 2x cash flow growth.

- Once net margin in traditional software/services expands by 2%-3%+ at the industry level thanks to AI agents, capital will rotate from hardware to applications, igniting a second major leg higher in AI software stocks.

2) Seamless mass production and "supply-constrained" status for next-gen flagship chips (e.g., NVIDIA (NVDA) Vera Rubin)

- NVIDIA (NVDA) plans to fully roll out its next-generation Vera Rubin architecture AI platform in 2H 2026.

- While Blackwell is already in steady production, if NVIDIA (NVDA) announces at upcoming tech summits that "Rubin will be supply-constrained throughout its lifecycle," it would directly dispel "2027 compute peak" fears, and leading-edge semis (including the SK Hynix (000660 KS) and Micron (MU) HBM4 chain) would keep climbing.

3) Sovereign AI — non-discretionary, nation-level spending

- Over the past two years, AI has been a capital arms race led by Microsoft (MSFT), Alphabet (GOOGL), and other U.S. "Magnificent Seven" tech giants. As countries recognize AI as a strategic asset tied to economic competitiveness, geopolitics, and national defense, "digital sovereignty" in 2026 formally evolves into a hard requirement for governments globally.

- The EU (rolling out the Tech Sovereignty Package), the Middle East (sovereign funds in Saudi Arabia and the UAE), Japan, India, and others will no longer rely solely on U.S. cloud hyperscalers; instead, they will fund—directly from national budgets—the build-out of "sovereign foundation models and national compute centers" fully native to local language, culture, and legal compliance.

- AI hardware and semiconductor buyers expand from "a few tech majors" to "finance ministries and sovereign funds across major countries." This nation-level arms race brings price-insensitive capital, pushing order visibility for cutting-edge chips (e.g., NVIDIA (NVDA), the TSMC (TSM US / 2330 TT) ecosystem), foundry, and advanced packaging out beyond 2030. Such rigid aggregate demand will collectively lift P/E multiples across the semiconductor supply chain.

7.4 Industry Downside Risks

Markets are highly fragile; any evidence of "circular financing" or "input–output imbalance" could trigger a plunge:

1) Commercial monetization underwhelms (ROI anxiety at the majors)

- This is the market's core concern today. Over the past two years, Microsoft (MSFT), Alphabet (GOOGL), Meta (META), Amazon (AMZN), and other Hyperscalers have invested USD hundreds of billions in CapEx, mainly on chips and data centers. Yet the conversion of these investments into actual revenue (subscriptions, cloud services growth) has been slow.

- Key risk: the market starts questioning AI ROI. If B2B or B2C software fails to produce killer apps commensurate with massive CapEx, slowing revenue growth could prompt Wall Street to mark down tech giants, triggering sector-wide Valuation De-rating.

- Underlying logic: The marginal cost of building and running LLMs is very high (inference costs, electricity). If ARPU cannot rise, elevated operating costs will severely compress margins.

2) Tech innovation leading to invalidation of legacy hardware investments (Technological Obsolescence)

- This is one of the largest, most opaque, and most vexing structural risks in the AI industry. In memory (e.g., HBM, high-end DRAM) and compute chips (e.g., GPU, ASIC), this risk is evolving at an unprecedented "brutal speed."

- In traditional manufacturing, machines can be depreciated over 10 years, but AI hardware—especially compute and high-bandwidth memory—now has a 2–3 year lifecycle. Take memory (HBM) as an example: two years ago the market was chasing HBM3, then quickly shifted to HBM3E; today (2026), HBM4 from SK Hynix (000660 KS), Micron (MU), and Samsung Electronics (005930 KS) is already dominating and tightly tied to NVIDIA (NVDA)'s next-gen massive compute platform.

- The cost of investment invalidation: Hyperscalers or compute lessors such as CoreWeave who spent billions of dollars a year ago on then-cutting-edge HBM3 servers will see rental rates collapse when far more power-efficient, multiples-faster HBM4/HBM4E ship at scale. These machines aren't physically broken, but economically they've become "stranded assets" or "stranded investments."

7.5 What if the industry sees a correction?

If the industry corrects (correction), we recommend more defensive sectors: Consumer staples and Pharma names with unusually cheap valuations and upward-revised EPS estimates.

9. KGI AI-30 Basket

9.1 Industry Valuation Comp Sheet

KGI Global AI-30 Basket - Valuation Comp Sheet

Company regional breakdown: US: 19; China: 5; Taiwan: 2; South Korea: 2; Netherlands: 1; Japan: 1

Ticker	Name	名稱	Global AI Layer	Country	Market Cap (USD mil)	Price Changes (%) 1M	3M	YTD	1YR	Last Price	BBG Tgt Price	Upside	Sales YoY (%) 2026	2027	EPS YoY (%) 2026	2027	Gross Margin 2026	2027	EBIT Margin 2026	2027	Net Margin 2026	2027	ROE 2026	2027	FCF Yield	EV/Sales	P/E 2026	2027	P/B 2026	2027
(1) Global AI Infra Monoply																														
NVDA US Equity	Nvidia	英偉達	6) Logic Chips	US	4,736,908	-8.9	14.3	5.0	26.3	196	301	53.9%	65.5	81.5	**59.5**	**88.7**	71.3	74.9	63.4	66.2	53.7	55.8	92.9	92.6	**2.5**	18.3	41.8	**21.8**	31.1	15.8
TSM US Equity	TSMC	台積電	4) Semiconductor Fabs/ Foundries	TW	2,256,064	5.5	33.4	43.1	94.2	435	469	7.9%	36.5	27.8	**49.2**	**26.7**	65.7	65.2	57.3	56.9	49.1	48.8	38.8	37.1	**1.9**	14.3	28.1	**22.1**	9.5	7.3
AVGO US Equity	Broadcom	博通	8) Networking & Interconnects	US	1,802,695	-10.2	22.5	9.5	40.2	379	523	38.0%	65.6	63.8	**69.4**	**66.2**	75.6	73.0	65.5	65.0	53.8	54.2	54.6	59.6	**1.8**	24.5	32.8	**19.7**	16.4	10.0
MU US Equity	Micron	美光科技	7) Memory	US	1,370,586	35.5	241.4	325.2	863.1	1,214	1,488	22.6%	238.2	89.8	**764.9**	**107.1**	80.3	85.2	74.6	81.6	64.1	69.1	74.5	71.8	**1.9**	14.9	16.9	**8.2**	10.6	5.0
005930 KS Equity	Samsung Electronics	三星電子	7) Memory	KR	1,251,908	9.8	84.2	175.7	451.5	332,000	465,649	40.3%	106.6	27.0	**495.6**	**38.5**	69.4	73.6	51.5	57.1	41.1	44.5	49.8	44.1	**2.9**	4.7	7.5	**5.4**	3.2	2.1
000660 KS Equity	SK Hynix	SK海力士	7) Memory	KR	1,241,648	29.0	188.4	313.7	820.5	2,697,000	3,109,772	15.3%	251.9	47.4	**394.9**	**44.8**	83.6	82.9	76.5	77.8	63.9	62.5	92.8	62.7	**2.2**	14.3	8.8	**6.1**	5.7	3.0
ASML NA Equity	ASML	艾司摩爾	3) Chip Manufacturing Equipment	NL	702,852	14.7	36.7	73.0	135.3	1,594	1,600	0.3%	19.6	23.2	**28.2**	**33.5**	52.5	54.1	36.6	40.1	31.5	33.8	55.5	58.9	**1.4**	18.2	50.3	**37.7**	25.7	19.0
				Simple Average		**10.8**	**88.7**	**135.0**	**347.3**			**25.5%**	**112.0**	**51.5**	**265.9**	**57.9**	**71.2**	**72.7**	**60.8**	**63.5**	**51.0**	**52.7**	**65.6**	**61.0**	**2.1**	**15.6**	**26.6**	**17.3**	**14.6**	**8.9**
				Market Cap Weighted Average		**4.4**	**65.8**	**93.1**	**244.8**			**33.1%**	**97.1**	**59.5**	**201.8**	**65.4**	**71.8**	**73.6**	**62.5**	**64.9**	**52.6**	**54.2**	**70.7**	**67.6**	**2.2**	**16.4**	**29.9**	**18.0**	**18.1**	**10.2**
(2) Mega 7																														
NVDA US Equity	Nvidia	輝達	6) Logic Chips	US	4,736,908	-8.9	14.3	5.0	26.3	196	301	53.9%	65.5	81.5	**59.5**	**88.7**	71.3	74.9	63.4	66.2	53.7	55.8	92.9	92.6	**2.5**	18.3	41.8	**21.8**	31.1	15.8
GOOGL US Equity	Alphabet	谷歌母公司	12) Hyperscale Cloud & AI Platforms	US	4,175,544	-11.6	22.4	9.8	98.1	344	432	25.7%	23.2	21.2	**11.1**	**7.4**	70.3	69.6	40.3	40.6	41.3	35.7	49.5	38.0	**1.5**	9.8	24.1	**22.4**	7.3	5.6
AAPL US Equity	Apple	蘋果	17) Horizontal Enterprise AI	US	4,041,226	-10.8	8.8	1.2	36.9	275	320	16.3%	14.9	9.0	**17.5**	**10.2**	48.2	47.8	32.4	31.9	26.9	26.8	132.5	108.0	**3.2**	8.8	31.4	**28.5**	35.4	24.8
MSFT US Equity	Microsoft	微軟	12) Hyperscale Cloud & AI Platforms	US	2,620,975	-15.2	-3.6	-27.0	-29.1	353	561	59.0%	16.9	16.7	**25.5**	**13.9**	67.8	66.6	46.6	46.7	38.8	37.7	31.1	27.6	**2.8**	8.4	20.6	**18.1**	5.9	4.6
AMZN US Equity	Amazon	亞馬遜	12) Hyperscale Cloud & AI Platforms	US	2,441,971	-14.4	9.4	-1.7	4.6	227	315	38.8%	15.1	13.2	**14.9**	**12.5**	51.2	52.6	12.6	14.1	13.8	13.8	19.8	17.8	**-0.1**	3.4	22.0	**19.6**	4.5	3.6
TSLA US Equity	Tesla	特斯拉	18) Vertical / Industry AI	US	1,408,847	-13.5	0.8	-16.6	15.1	375	422	12.6%	8.2	15.2	**12.8**	**30.0**	19.6	20.1	5.3	6.7	6.4	7.3	6.4	7.6	**0.5**	14.1	200.3	**154.1**	14.4	13.2
META US Equity	Meta	Meta平台	12) Hyperscale Cloud & AI Platforms	US	1,378,034	-11.3	-0.9	-17.8	-25.2	543	820	51.0%	25.9	19.2	**27.9**	**-4.1**	80.3	79.1	35.7	34.9	40.9	34.0	32.1	25.5	**3.5**	6.4	13.5	**14.1**	4.5	3.6
				Simple Average		**-12.2**	**7.3**	**-6.7**	**18.1**			**36.8%**	**24.2**	**25.2**	**24.2**	**22.7**	**58.4**	**58.7**	**33.7**	**34.4**	**31.7**	**30.2**	**52.0**	**45.3**	**2.0**	**9.9**	**50.5**	**39.8**	**14.7**	**10.2**
				Market Cap Weighted Average		**-11.7**	**10.1**	**-2.6**	**29.1**			**36.8%**	**28.6**	**30.5**	**26.9**	**28.6**	**60.9**	**61.5**	**38.9**	**39.7**	**35.4**	**34.2**	**65.6**	**57.5**	**2.1**	**10.7**	**40.1**	**30.9**	**18.0**	**11.7**
(3) Segment Leaders																														
AMD US Equity	Advanced Micro Devices	超微半導體	6) Logic Chips	US	868,409	5.7	161.4	148.7	270.7	533	498	-6.4%	43.7	53.8	**73.8**	**81.0**	55.5	55.6	26.8	31.7	24.2	28.9	14.2	21.5	**1.0**	23.0	73.5	**40.6**	12.2	10.3
6723 JP Equity	Renesas Electronics	瑞薩電子	6) Logic Chips	JP	59,093	11.4	113.9	138.8	194.4	5,110	4,796	-6.1%	16.9	9.8	**37.3**	**5.6**	58.0	58.8	24.7	23.0	30.0	27.4	14.6	12.2	**3.7**	7.5	20.4	**19.3**	3.3	3.0
SNDK US Equity	Sandisk	晟碟	7) Memory	US	345,790	46.9	287.1	883.7	4,822	2,335	1,845	-21.0%	169.0	130.5	**2,095.5**	**191.3**	69.7	82.6	59.9	78.4	51.6	66.1	60.5	71.0	**1.3**	26.0	35.6	**12.2**	18.6	7.3
MRVL US Equity	Marvell	邁威爾	8) Networking & Interconnects	US	246,046	35.1	187.9	231.0	251.7	281	256	-9.0%	42.1	40.6	**80.9**	**42.4**	59.5	58.6	35.3	36.3	30.1	31.7	2.8	12.0	**0.7**	28.4	99.1	**69.6**	16.8	13.8
COHR US Equity	Coherent	連貫	8) Networking & Interconnects	US	79,674	6.8	67.4	120.6	353.1	407	386	-5.3%	21.5	36.2	**54.5**	**52.6**	39.4	40.5	20.4	22.4	15.1	17.2	9.9	12.9	**-0.7**	12.3	74.7	**48.9**	7.6	6.9
LITE US Equity	Lumentum	魯門特姆	8) Networking & Interconnects	US	67,061	-5.4	25.1	133.9	810.1	862	1,126	30.7%	82.2	85.7	**295.9**	**121.6**	45.3	49.7	29.1	39.0	25.1	33.1	32.7	44.8	**0.2**	27.0	105.7	**47.7**	18.9	12.0
GEV US Equity	GE Vernova	奇異維諾瓦	9) Power Delivery & Power Equipment	US	291,687	1.4	24.3	66.1	114.2	1,085	1,219	12.3%	19.4	14.3	**142.0**	**16.0**	23.6	26.8	10.2	14.7	13.9	12.6	39.8	36.1	**2.6**	7.3	52.7	**45.4**	18.7	14.5
ETN US Equity	Eaton	伊頓	9) Power Delivery & Power Equipment	US	163,036	4.2	17.6	31.8	20.6	420	460	9.5%	16.3	10.5	**10.1**	**16.8**	37.6	38.7	19.8	21.0	16.0	17.1	23.5	24.5	**2.3**	6.5	31.6	**27.0**	7.3	6.4
VRT US Equity	Vertiv	維諦	10) Cooling Systems	US	125,054	0.5	29.0	101.0	163.0	326	381	17.0%	35.6	28.2	**54.0**	**34.0**	38.5	39.3	23.2	24.5	18.0	18.8	48.5	46.2	**1.8**	11.6	50.3	**37.6**	21.2	14.8
3017 TT Equity	Asia Vital Components	奇鋐科技	10) Cooling Systems	TW	28,337	-15.6	1.1	52.3	212.5	2,300	3,360	46.1%	57.8	35.5	**92.3**	**44.3**	30.4	31.3	25.0	26.3	16.9	18.0	58.1	57.1	**4.2**	5.3	24.3	**16.9**	12.3	8.6
				Simple Average		**9.1**	**91.5**	**190.8**	**721.2**			**6.8%**	**50.4**	**44.5**	**293.6**	**60.6**	**45.7**	**48.2**	**27.4**	**31.7**	**24.1**	**27.1**	**30.4**	**33.8**	**1.7**	**15.5**	**56.8**	**36.5**	**13.7**	**9.8**
				Market Cap Weighted Average		**13.8**	**137.7**	**244.9**	**932.8**			**-2.3%**	**56.9**	**53.4**	**390.2**	**75.9**	**50.7**	**53.3**	**29.7**	**35.6**	**26.5**	**30.9**	**26.8**	**32.0**	**1.4**	**19.3**	**62.6**	**38.5**	**14.5**	**10.5**
(4) Strategic Reshaping																														
Intc US Equity	Intel	英特爾	4) Semiconductor Fabs/Foundries	US	667,805	7.6	201.3	260.1	490.5	133	100	-25.1%	11.0	11.5	**157.6**	**43.7**	40.4	42.9	12.2	15.4	9.6	12.6	3.6	6.5	**-0.5**	12.7	122.8	**85.4**	5.6	5.2
Dell US Equity	Dell	戴爾	11) Data Center Construction & Operation	US	264,563	34.2	132.9	225.3	225.0	409	486	18.6%	18.8	51.5	**26.5**	**74.4**	20.4	17.8	8.7	9.3	6.1	7.0	-353.7	-422.9	**3.6**	2.1	41.1	**22.8**	N/A N/A	229.3
				Simple Average		**20.9**	**167.1**	**242.7**	**357.8**			**-3.2%**	**14.9**	**31.5**	**92.1**	**59.1**	**30.4**	**30.3**	**10.4**	**12.3**	**7.9**	**9.8**	**-175.1**	**-208.2**	**1.6**	**7.4**	**82.0**	**54.1**	**5.6**	**117.2**
				Market Cap Weighted Average		**15.1**	**181.9**	**250.2**	**415.2**			**-12.7%**	**13.2**	**22.8**	**120.4**	**52.4**	**34.7**	**35.8**	**11.2**	**13.6**	**8.6**	**11.0**	**-97.8**	**-115.4**	**0.7**	**9.7**	**99.6**	**67.7**	**4.0**	**68.8**
(5) China AI																														
700 HK Equity	Tencent	騰訊	17) Horizontal Enterprise AI	CN	478,767	-6.1	-16.8	-31.2	-19.6	412	687	66.7%	10.3	9.4	**8.3**	**11.0**	56.8	57.4	32.6	33.2	33.3	33.5	19.3	18.6	**6.2**	4.3	11.9	**10.7**	2.4	2.1
BABA US Equity	Alibaba	阿里巴巴	12) Hyperscale Cloud & AI Platforms	CN	228,240	-26.6	-24.2	-35.1	-16.6	95	186	95.2%	2.7	9.6	**-59.0**	**59.4**	40.4	39.5	6.3	8.1	7.5	9.5	8.0	8.7	**-3.2**	1.1	20.1	**15.1**	1.6	1.5
2513 HK Equity	Knowledge Atlas (Zhipu)	智譜	14) Foundation Models	CN	125,080	63.6	221.6	N/A	N/A	2,200	1,373	-37.6%	346.2	173.0	**-32.8**	**6.2**	34.9	36.8	-147.0	-55.6	-135.6	-48.6	46.3	554.2	**N/A**	1,172.7	N/A	**N/A**	N/A	N/A
2476 HK Equity	Victory Giant	勝宏科技	8) Networking & Interconnects	CN	46,430	-23.2	N/A	N/A	N/A	336	540	60.7%	75.1	70.6	**95.8**	**77.6**	40.0	41.4	31.0	32.3	26.8	28.1	30.8	38.1	**-1.0**	15.8	33.2	**18.7**	9.4	6.8
BIDU US Equity	Baidu	百度	6) Logic Chips	CN	35,383	-18.0	-5.4	-20.4	20.3	104	177	70.2%	2.7	7.7	**-5.6**	**18.1**	40.9	42.4	10.2	11.5	13.8	14.6	5.3	5.9	**-4.1**	1.0	14.0	**11.9**	0.8	0.8
				Simple Average		**-2.1**	**43.8**	**-28.9**	**-5.3**			**51.0%**	**87.4**	**54.1**	**1.3**	**34.4**	**42.6**	**43.5**	**-13.4**	**5.9**	**-10.8**	**7.4**	**21.9**	**125.1**	**-0.5**	**239.0**	**19.8**	**14.1**	**3.6**	**2.8**
				Market Cap Weighted Average		**-3.0**	**15.3**	**-25.9**	**-13.6**			**59.4%**	**57.4**	**34.9**	**-10.3**	**26.1**	**48.2**	**48.7**	**0.5**	**13.9**	**2.6**	**15.3**	**20.2**	**89.9**	**2.2**	**163.9**	**13.5**	**10.8**	**2.2**	**1.8**
				KGI AI 30 - Simple Average		**4.0**	**71.8**	**114.3**	**373.0**			**23.3%**	**62.0**	**41.1**	**169.9**	**45.7**	**52.2**	**53.5**	**27.6**	**33.0**	**24.3**	**28.4**	**26.5**	**39.9**	**1.5**	**50.8**	**45.1**	**31.3**	**11.8**	**15.9**
				KGI AI 30 - Market Cap Weighted Average		**-3.0**	**45.3**	**57.6**	**186.2**			**29.5%**	**53.0**	**36.3**	**118.2**	**38.6**	**62.0**	**62.8**	**42.4**	**44.3**	**37.4**	**37.7**	**55.4**	**50.9**	**2.0**	**16.6**	**38.2**	**28.1**	**15.1**	**11.7**

Source: Bloomberg, KGI (As of 26/06/2026)

9.2 Company Pages: Fundamental/Valuation/Catalysts/Risks

US equities: NVIDIA (NVDA)

1. Investment Thesis (Investment Thesis)
Q1 revenue reached USD 81.6 billion, marking 14 consecutive quarters beating guidance; Data Center revenue surged to USD 75.2 billion. Non-hyperscale (ACIE) and Sovereign AI revenue grew 31% QoQ to a 50% mix, underscoring exceptionally strong non-hyperscale demand. As AI shifts from single-pass inference to Reasoning and Agentic AI, revenue across the accelerated computing stack should further accelerate, with leadership difficult to challenge.

2. Valuation Summary (Valuation Summary)
Target price raised to USD 301. Beyond maintaining cumulative USD 1 trillion revenue expectations for Blackwell and Rubin, the Vera CPU designed for Agentic AI unlocks a new USD 200 billion market; this year, standalone CPU revenue visibility is nearly USD 20 billion, providing strong rerating support and premium optionality.

3. Share Price Catalysts (Share Price Catalysts)

Key catalysts include: next-gen Vera Rubin platform starts shipping and ramps from CY3Q26, with inference throughput 35x Blackwell; 2Q27 revenue guidance at USD 91 billion with potential to set new highs; Spectrum-X networking revenue far exceeds all Ethernet peers combined; and Embodied AI momentum is recovering, with the driverless taxi fleet in partnership with Uber expanding to nearly 30 cities in 2028.

4. Investment Risks (Investment Risks)
Main risks include: near-term supply chain facing HBM4 supply shortages, and Rubin's thermal redesign pushing volume shipments to 4Q and 1Q next year; secondly, increased R&D investment and AI tools usage leading the company to raise full-year opex growth to as high as 40%; finally, intensifying competition in chips and macro volatility potentially suppressing consumer-edge compute demand.

5. Company Profile (Company Profile)
The company is a global pioneer in GPUs and accelerated computing, offering a full-stack chips-and-software system solution comprising GPU, Vera CPU, and DPU. NVIDIA (NVDA), leveraging its dominant share in AI training and inference, is transforming traditional data centers into high-performance "AI factories," while aggressively expanding in Agentic AI, autonomous driving, and embodied robots, leading the global AI technology transformation.

Financials (USD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	130.5	215.9	392.0	558.1	681.3
Adj. EBIT (100 Mn)	81.5	130.4	260.4	372.4	442.1
Adj. EPS (CNY)	3.0	4.8	9.0	12.8	14.7
OCF (100 Mn)	64.1	102.7	215.2	308.5	378.9
FCF (100 Mn)	60.9	96.7	205.9	292.1	360.9
Margins & Growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue Growth	114.2%	65.5%	81.5%	42.4%	22.1%
Adj. EPS Growth	150.6%	59.5%	88.7%	42.3%	14.5%
Gross Margin	75.3%	71.2%	74.8%	74.3%	72.9%
Adj. EBIT Margin	62.4%	60.4%	66.4%	66.7%	64.9%
Net Margin	55.8%	55.6%	55.9%	55.5%	53.5%
Net Debt/Equity	-41.5%	-32.5%	-96.1%	-204.6%	-339.3%
Return on Capital	119.2%	101.5%	92.0%	77.2%	60.9%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF Yield (%)	1.2%	1.9%	4.0%	5.7%	7.1%
Dividend Yield (%)	0.1%	0.1%	0.4%	0.4%	0.9%
EV/EBITDA(x)	41.0	33.6	18.8	13.2	10.6
P/E (x)	48.5	38.3	23.4	16.4	14.4
P/B (x)	44.0	29.0	17.0	10.7	7.1
P/S (x)	26.8	21.2	13.0	9.1	7.5

Volatility

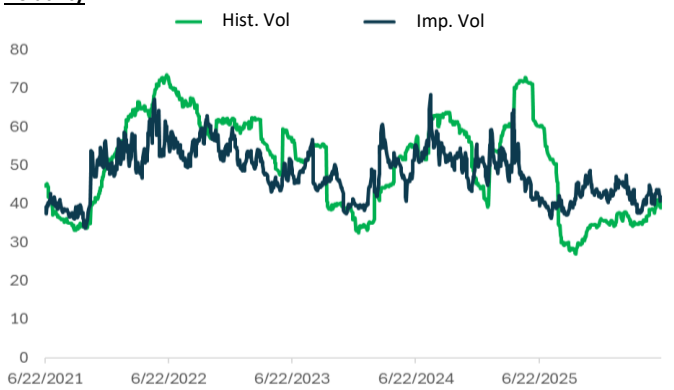

Valuation

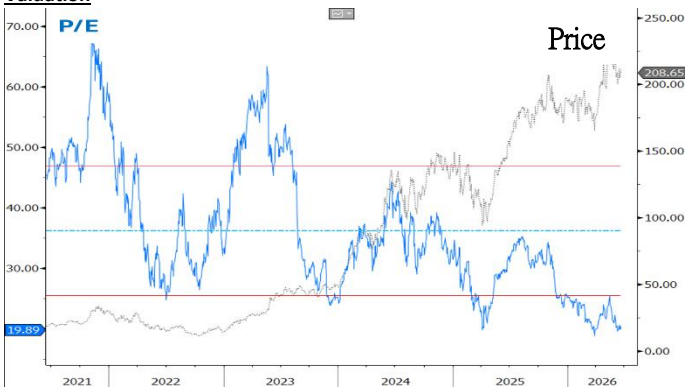

Source of revenue

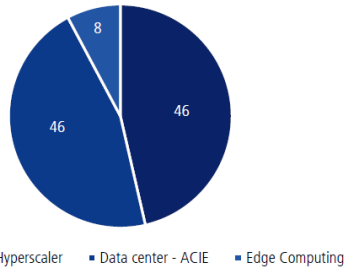

Taiwan Equities: TSMC (2330.TW)

1. Investment Thesis (Investment Thesis)

Benefiting from the global AI demand surge, TSMC is delivering strong earnings and expansion momentum. 1Q26 results were robust: quarterly revenue reached TWD 1.13 trillion (+35.1% y/y), EPS TWD 22.08, GM 66.2% above expectations. The company raised 2026 full-year revenue growth outlook to >30% y/y and guided Capex skewing toward the high end at USD 52–56 billion. Core N3 is supply-constrained; TSMC announced capacity expansions in Tainan, Arizona, and Kumamoto, expected to lift N3 capacity by 40–45% by end-2028, with N3 GM above the corporate average starting 2H26. CoWoS advanced packaging capacity is inflecting, with monthly capacity projected to rise from 72,000 wafers by end-2025 to 120,000 and 170,000 by end-2026 and end-2027, respectively. The next-gen A14 node is on track, set for volume production in 2028, introducing 2nd-gen nanosheet architecture; at iso-power, performance improves 10–15%, cementing its absolute leadership in AI foundry.

2. Valuation Summary (Valuation Summary)

Current valuation is attractive; Target Price raised to TWD 2,756. Based on 22x 2027E P/E, in line with historical upcycle averages. Given better-than-expected process cost improvements and profit mix optimization, 2026–27 EPS estimates are raised by 6% and 10%, respectively; high visibility in advanced nodes supports continued multiple expansion and re-rating.

3. Share Price Catalysts (Share Price Catalysts)

Key catalysts include: (1) N3 supply-demand tightness, with announced fab expansions in Tainan, the US, and Japan; GM to exceed the company average from 2H26; (2) 2Q26 revenue guidance up 10% q/q, well above market expectations; (3) high-margin advanced packaging CoWoS capacity doubling, and next-gen A14 (1.4nm) progressing smoothly, with mass production expected in 2028 as scheduled.

4. Investment Risks (Investment Risks)

Potential risks include: first, heightened Middle East and geopolitical tensions may push up power and key chemical raw material costs; while the company has built safety stock and diversified suppliers, cost structure volatility warrants monitoring; second, global macro uncertainty could dampen the recovery pace of smartphones and other consumer electronics end-demand; finally, aggressive industry expansion in advanced nodes and packaging capacity could lead to future oversupply.

5. Company Profile (Company Profile)

The world's largest pure-play semiconductor foundry with leading global market share. TSMC provides wafer fabrication, testing, advanced packaging, and design services, and continues to expand new fabs across Taiwan, the US, and Japan. With absolute leadership in 3nm, 5nm, and future A14 technologies, the company has become the foundation for global HPC, smartphones, and AI chips.

Financials (TWD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	2894.3	3809.1	5196.1	6600.9	8080.1
Adj. EBIT (100 Mn)	1322.0	1936.1	2984.0	3766.0	4559.5
Adj. EPS (TWD)	45.3	66.3	98.7	124.3	152.0
OCF (100 Mn)	1887.4	2358.7	3208.8	4087.0	4842.1
FCF (100 Mn)	931.4	1086.3	1485.8	2039.5	2675.3
Margins & Growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue Growth	33.9%	31.6%	36.4%	27.0%	22.4%
Adj. EPS Growth	39.9%	46.4%	49.0%	25.9%	22.3%
Gross Margin	56.1%	59.9%	65.7%	65.2%	64.8%
Adj. EBIT margin	45.7%	50.8%	57.4%	57.1%	56.4%
Net margin	40.5%	45.1%	48.9%	48.4%	48.2%
Net debt/equity	-31.7%	-36.7%	-52.9%	-77.4%	-112.8%
Return on capital	30.3%	35.4%	38.0%	36.2%	30.5%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	1.4%	1.7%	2.3%	3.1%	4.1%
Dividend yield (%)	0.2%	0.2%	1.0%	1.2%	1.4%
EV/EBITDA(x)	13.4	14.6	16.6	13.2	10.8
P/E (x)	23.8	23.4	25.4	20.2	16.5
P/B (x)	6.5	7.4	8.6	6.6	4.9
P/S (x)	9.6	10.6	12.5	9.9	8.1

Volatility

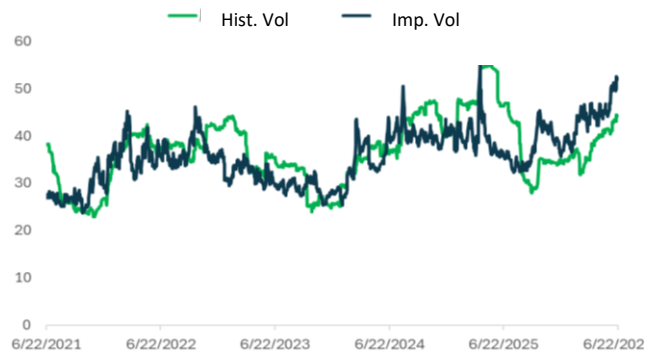

Valuation

Revenue breakdown by node

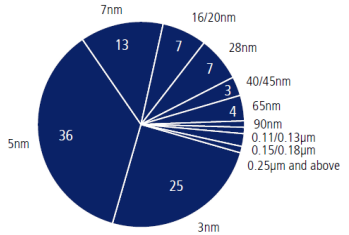

US: Broadcom (AVGO)

1. Investment Thesis

Broadcom holds dominant leadership in custom silicon (ASIC/XPU) and high-end networking. Latest quarterly revenue reached USD 22.187 billion, with AI revenue up 145% YoY to USD 10.8 billion; AI bookings exceeded USD 30 billion. Next quarter revenue is guided to surge to USD 29.4 billion, with AI at USD 16.0 billion (+50% QoQ). Driven by large custom chip orders from Google, Anthropic, and OpenAI, FY2026–FY2028 EPS expected to grow at high rate. In addition, the sharp increase in server core counts is accelerating growth in the higher-margin VMware segment (quarterly guide USD 8.9 billion, +31% YoY). Dual-engine growth, coupled with a 67.3% OM (+200 bps YoY), will significantly enhance operating leverage.

2. Valuation Summary

Market target price range has been raised to USD 522. Current P/E is ~22.7x, near the low end of its 20–40x historical range, offering attractive and defensive valuation. Most analyses apply SOTP; based on FY2027–FY2028 EPS estimates (USD 19.2 and USD 25.8), a reasonable 27x P/E supports higher target prices.

3. Share Price Catalysts

Catalysts include volume ramp from top customers' chips and capacity expansion. Meta will deliver the first 1 GW order in 2H27; custom chips for OpenAI and Anthropic enter mass production in 2027–28. In addition, the Singapore advanced packaging plant will start production ahead of schedule in August this year, easing capacity bottlenecks; accelerated adoption of next-gen networking silicon Tomahawk 6, and the surge in server core counts driving VMware ARR growth, should keep earnings revisions trending higher.

4. Investment Risks

Key risks center on high customer concentration, design/development complexity, and power infrastructure delays. The primary risk is that roughly 35–40% of revenue depends on the largest customer, Alphabet (GOOGL). In addition, ASIC chip design is highly complex; slippage in development timelines or market share versus expectations could create noise. Moreover, global data centers face delays in Power Shell power-supply buildouts; if customers cannot build out infrastructure on schedule, deployments will be delayed even if chips ship on time.

5. Company Profile

Broadcom (AVGO) is a global leader in semiconductors and infrastructure software. On semiconductors, it holds a dominant edge in high-speed Ethernet switching chips (Tomahawk, Jericho series) and custom AI ASICs. On software, Broadcom has expanded into high-margin enterprise virtualization and cloud software via a series of strategic M&A. The company now runs a dual-engine model across semiconductors and infrastructure software with high operating leverage.

Financials (USD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	51.6	63.9	105.8	173.1	226.1
Adj. EBIT (100 Mn)	13.5	25.5	70.1	113.9	151.1
Adj. EPS (USD)	4.9	6.8	11.6	19.2	25.8
OCF (100 Mn)	20.0	27.5	52.6	92.8	127.2
FCF (100 Mn)	19.4	26.9	51.1	90.7	118.1
Margins and Growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue Growth	44.0%	23.9%	65.6%	63.6%	30.7%
Adj. EPS Growth	47.7%	40.0%	69.6%	66.4%	34.2%
Gross Margin	63.5%	67.9%	75.0%	73.1%	72.4%
Adj. EBIT Margin	26.1%	39.9%	66.3%	65.8%	66.8%
Net Margin	11.4%	36.2%	54.0%	54.7%	54.8%
Net Debt/Equity	88.0%	61.9%	26.2%	-62.0%	-166.7%
Return on Capital	12.9%	31.0%	56.8%	62.5%	55.4%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF Yield (%)	1.0%	1.4%	2.6%	4.6%	6.0%
Dividend Yield (%)	0.2%	0.2%	0.7%	0.9%	1.1%
EV/EBITDA(x)	32.8	51.3	27.8	17.1	13.0
P/E (x)	137.3	77.5	35.6	21.4	15.9
P/B (x)	11.7	21.6	17.8	10.9	7.2
P/S (x)	15.1	27.3	18.5	11.3	8.7

Volatility

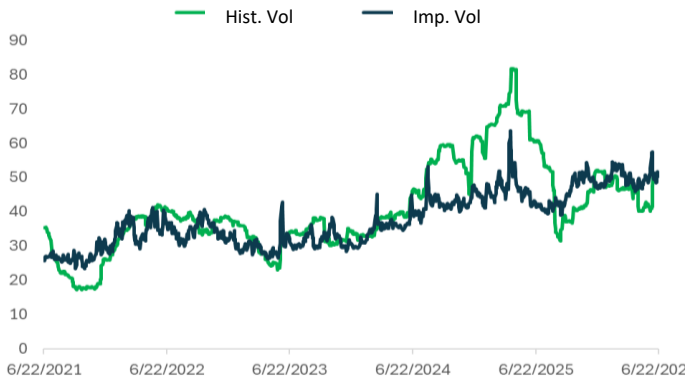

Valuation

Segment Revenue Share

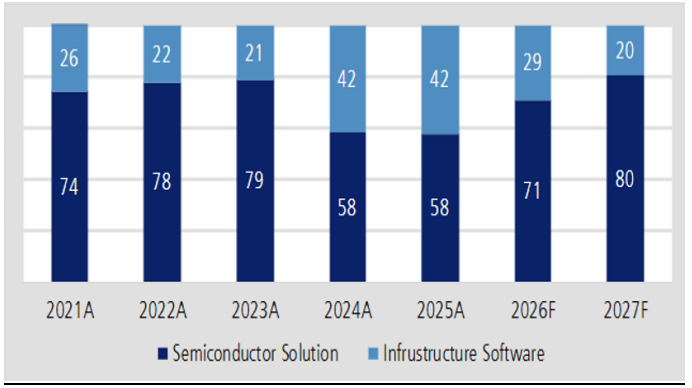

Korea: SK Hynix (000660 KS)

1. Investment Thesis (Investment Thesis)
Benefiting from an AI demand boom, the core business shows exceptional resilience. We forecast 2026 DRAM and NAND ASPs to jump 200% and 186%, with server DDR5 and SSD ASPs soaring 329% and 267%. HBM's position remains solid, with shipment share reaching 50% in 2026. The first batch of 3-year LTA signings locks in 25% of DRAM and 20% of NAND volumes, shifting the earnings model from cyclical to structural. This underpins distributing KRW 240 trillion in FCF over 25A-27E, driving a rerating from P/B to P/E.

2. Valuation Summary (Valuation Summary)
The valuation framework is shifting from P/B to P/E to reflect earnings visibility from LTAs. We adopt SOTP, assigning 9.5x EV/EBITDA to the highly customized HBM business (akin to a foundry model) and 7.7x to commodity memory. We raise TP to KRW 3,100,000, implying only 8x 2026E P/E and substantial rerating upside.

3. Share Price Catalysts (Share Price Catalysts)
Key catalysts include: (1) overseas ADR issuance followed by share buybacks, directly supporting the share price; (2) disclosure of 3-year LTA negotiation details, materially enhancing earnings visibility; (3) highly attractive shareholder returns, with up to KRW 240 trillion FCF to be distributed in 2025–2027 (including dividends and share cancellations); (4) 2027 HBM4 pricing upside surprise.

4. Investment Risks (Investment Risks)
Key downside risks include: (1) LTA negotiations progressing slower than expected, delaying rerating; (2) weak macro or geopolitical tensions prompting CSPs to slow AI Capex, suppressing memory demand; (3) innovations reducing DRAM required per AI rack; (4) industry capacity expansion faster than expected, flipping supply-demand.

5. Company Profile (Company Profile)
The company is a pure-play global memory semiconductor leader focused on R&D and manufacturing of DRAM and NAND flash. By revenue, it is the world's No. 2 in DRAM and No. 3 in NAND. It holds a strong leadership position in AI and is the global leader in High Bandwidth Memory (HBM). The company joined the SK Group in 2012 and has major production bases in Korea and China.

Financials (KRW)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (Bn)	66,193	97,147	341,261	497,858	546,042
Adj. EBIT (Bn)	23,467	47,206	262,437	388,428	410,186
Adj. EPS (KRW)	27,182	62,161	304,220	436,588	469,831
OCF (Bn)	29,796	53,373	178,825	284,627	316,948
FCF (Bn)	13,850	25,854	144,600	233,880	264,856
Margins and growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue growth	102.0%	46.8%	251.3%	45.9%	9.7%
Adj. EPS growth	-305.2%	128.7%	389.4%	43.5%	7.6%
Gross margin	48.1%	60.4%	83.5%	82.7%	79.8%
Adj. EBIT margin	35.5%	48.6%	76.9%	78.0%	75.1%
Net margin	29.9%	44.2%	61.7%	60.3%	59.3%
Net debt/equity	15.3%	-8.4%	-115.9%	-293.7%	-517.3%
Return on capital	31.1%	44.2%	91.2%	60.8%	41.1%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	0.7%	1.3%	7.0%	11.4%	12.9%
Dividend yield (%)	0.0%	0.0%	0.4%	0.7%	0.9%
EV/EBITDA(x)	3.6	7.3	7.2	4.9	4.5
P/E (x)	6.1	10.8	9.5	6.6	6.1
P/B (x)	1.6	3.8	6.2	3.3	2.1
P/S (x)	1.8	4.6	6.0	4.1	3.8

Volatility

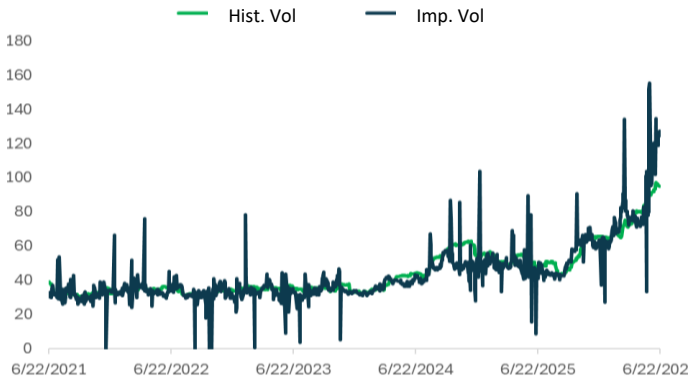

Valuation

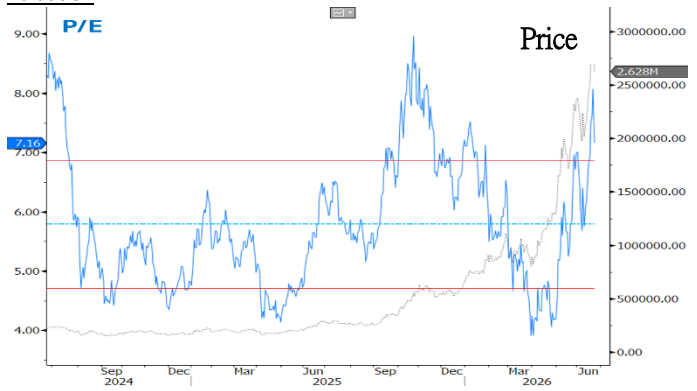

Source of Revenue

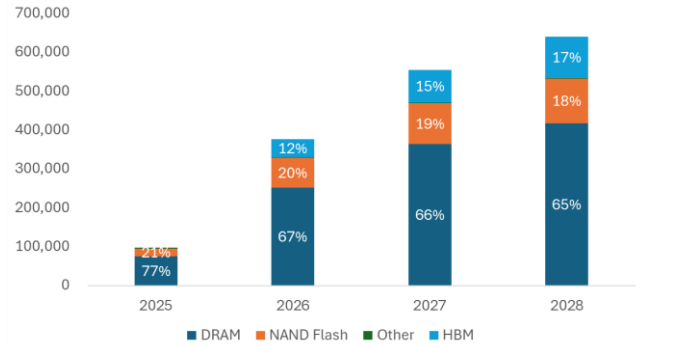

US Equities: Micron Technology (MU)

1. Investment Thesis

The AI boom is driving demand for high-capacity DRAM and enterprise SSDs, with query context window size up 30x YoY. Supply is structurally tight as HBM consumes more than 3x the wafers of standard DRAM. Industry expects 2026 global DRAM and NAND ASPs to surge 200% and 186%; since the start of 2026, spot prices are up 52% (premium contract prices +21%). Micron (MU) has strong pricing power; we expect FY2026 Q3 revenue to climb to USD 33.5 billion–USD 36.0 billion, with GM above 81%. More importantly, the company has shifted up to 30% of DDR shipments into multi-year Strategic Customer Agreements, fundamentally dampening traditional cyclicality and building a wide moat.

2. Valuation Summary

Street PTs have been raised to USD 1,507. Based on FY2027E EPS USD 117.2, the forward P/E is only ~12x, well below the 17x historical peak, offering both high growth and defensiveness. As HBM3E/HBM4 ramp and enterprise SSD share rises, blended GM is expected to climb to ~90% by end-2027; capacity locked by Strategic Customer Agreements (SCA) provides a high-certainty earnings floor for the mid-cycle.

3. Share Price Catalysts

The most immediate catalyst is the upcoming Q3 print and outlook, which we expect to beat. On the technology front, cost-advantaged 1-gamma DRAM and G9 NAND will become shipment mainstays by mid-2026, and next-gen HBM4 should ramp in 2027. Micron (MU)'s first AI-optimized PCIe Gen 6 data center SSD is certified on NVIDIA (NVDA)'s new Rubin reference platform; combined with more hyperscalers signing multi-year strategic long-term deals, this should sustain upward share-price momentum.

4. Investment Risks

Three key risks: (1) Overexpansion of Capex. Micron (MU) lifted FY2026 Capex to USD 25.0 billion+, and a concentrated capacity release post-2028 could reprise oversupply. (2) Some next-gen GPU platforms (e.g., Rubin) may face advanced packaging bottlenecks, prompting short-term DRAM de-spec'ing and weighing on shipments. (3) A global macro downturn could slow memory inventory digestion in legacy consumer electronics (e.g., PCs, smartphones).

5. Company Profile

Micron Technology (MU) is the world's No.3 DRAM and No.4 NAND supplier, with a vertically integrated global footprint. Its customers include top-tier PC/server OEMs, hyperscalers, and smartphone vendors. The company leverages advantages in HBM, high-capacity DDR5, and high-performance enterprise SSDs.

Financial Data (USD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	25.1	37.4	114.4	198.8	211.7
Adj. EBIT (100 Mn)	1.3	9.8	82.5	160.2	163.8
Adj. EPS (USD)	1.3	8.3	62.1	117.2	121.7
OCF (100 Mn)	8.5	17.5	66.1	141.8	159.8
FCF (100 Mn)	0.1	1.7	39.1	102.8	124.5
Margins and Growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue Growth	61.6%	48.9%	206.2%	73.7%	6.5%
Adj. EPS Growth	-124.3%	537.7%	649.5%	88.7%	3.8%
Gross Margin	22.3%	39.8%	77.6%	81.9%	78.8%
Adj. EBIT Margin	5.2%	26.1%	72.1%	80.6%	77.4%
Net profit margin	3.1%	22.8%	61.6%	66.8%	63.5%
Net debt/equity	10.9%	6.3%	-53.6%	-185.4%	-342.6%
Return on capital	1.7%	17.2%	71.8%	65.1%	42.7%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	0.0%	0.1%	3.1%	8.0%	9.7%
Dividend yield (%)	0.0%	0.0%	0.1%	0.1%	0.1%
EV/EBITDA(x)	12.3	7.6	14.2	7.6	7.1
P/E (x)	136.5	16.1	18.3	9.7	9.3
P/B (x)	2.3	2.5	10.6	5.2	3.2
P/S (x)	4.2	3.6	11.2	6.4	6.0

Volatility

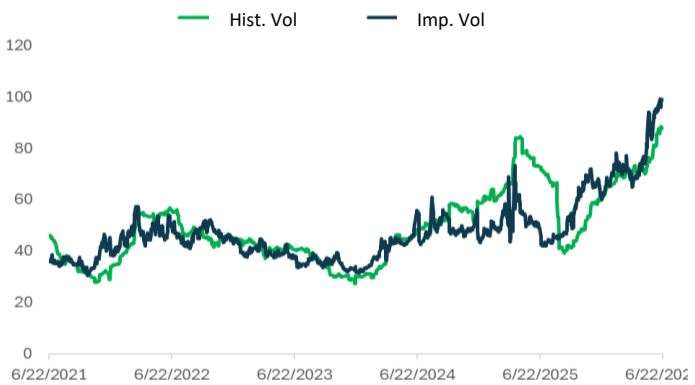

Valuation

Revenue

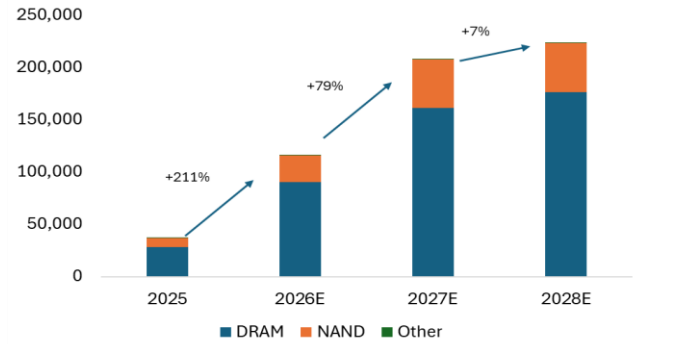

Korea: Samsung Electronics (005930 KS)

1. Investment Thesis (Investment Thesis)
Driven by the agentic AI boom, memory undersupply is likely to persist through 2028. We expect 3Q26 and 4Q26 contract prices to rise 17% qoq and 12% qoq, respectively, and up to 60% of server DDR5 to be locked in via LTAs. In HBM, 2027 bit shipments are estimated to reach 24 billion Gb, taking global share to 41%; in the same year, HBM4E and HBM4 contract pricing nears USD 3.5 and USD 3.0 per Gb, respectively, lifting overall ASP 67% YoY. Alongside the newly announced HBM5 combining HPB cooling and 2nm process, the technology moat will continue to underpin high growth.

2. Valuation Summary (Valuation Summary)
The stock currently trades around KRW 329,000–360,000, with a fair target range of KRW 460,000. We apply SOTP, assigning 7.9x to Memory, 4.1x to Foundry, and 4.8x to Mobile. With long-term ROE raised to 30.4%, the TP implies 10x P/E; current levels offer strong defensiveness and upside.

3. Share Price Catalysts (Share Price Catalysts)
Key catalysts: (1) labor-management settlement removing strike risk; (2) HBM5 with 2nm and HPB cooling unveiled at Computex, cementing tech leadership; (3) smooth progress on LTAs with major US CSPs; (4) the late-July quarterly results likely to signal further upside to memory contract prices and a shareholder return plan of 50% of FCF.

4. Investment Risks (Investment Risks)
Key risks to watch include: (1) HBM qualification at major customers progressing slower than expected; (2) weak PC and smartphone demand dragging on the recovery of commodity memory; (3) aggressive capacity expansion by peers leading to price corrections; (4) KRW appreciation eroding export profitability; (5) the new labor agreement (DS segment allocating 12% of operating profit as bonuses) will raise opex pressure in the near term.

5. Company Profile (Company Profile)
Founded in 1969, the company is a global leader in electronics and semiconductors. Operations are divided into two major segments: Device Solutions (DS), covering semiconductors and display panels, and Digital Media and Communications (DX), spanning mobile communications and consumer electronics. The company ranks No.1 globally in memory (approx. 45% share), OLED, and TVs, and No.2 by smartphone revenue, with a strong vertical-integration advantage.

Financials (KRW)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (Bn)	300,871	333,606	687,230	870,501	922,194
Adj. EBIT (Bn)	32,726	43,601	353,226	490,336	503,092
Adj. EPS (KRW)	5,632	7,467	44,284	60,973	62,873
OCF (Bn)	72,983	85,315	285,202	384,542	415,944
FCF (Bn)	21,576	37,793	212,516	330,349	358,083
Margins and growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue growth	16.2%	10.9%	106.0%	26.7%	5.9%
Adj. EPS growth	164.3%	32.6%	493.1%	37.7%	3.1%
Gross margin	38.0%	39.4%	69.3%	73.4%	68.8%
Adj. EBIT margin	10.9%	13.1%	51.4%	56.3%	54.6%
Net margin	11.2%	13.3%	40.8%	43.9%	42.3%
Net debt/equity	-23.2%	-23.1%	-63.8%	-125.3%	-197.3%
Return on capital	7.9%	9.5%	48.2%	42.8%	31.0%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	1.0%	1.8%	10.3%	16.0%	17.3%
Dividend yield (%)	0.1%	0.1%	1.4%	1.7%	1.9%
EV/EBITDA(x)	3.1	6.8	4.9	3.6	3.4
P/E (x)	10.7	18.2	8.0	5.8	5.6
P/B (x)	0.8	1.7	3.4	2.3	1.7
P/S (x)	1.1	2.1	3.0	2.4	2.2

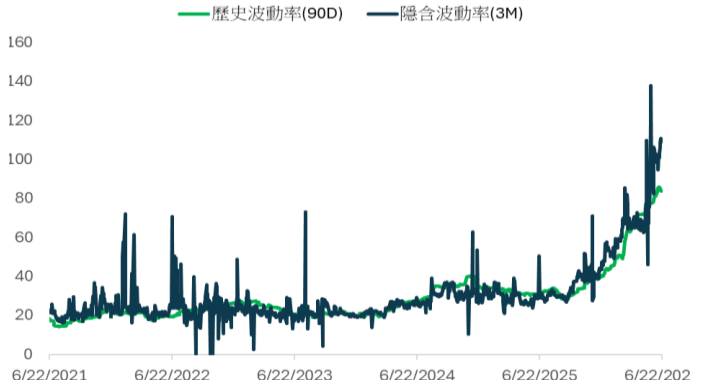

Valuation

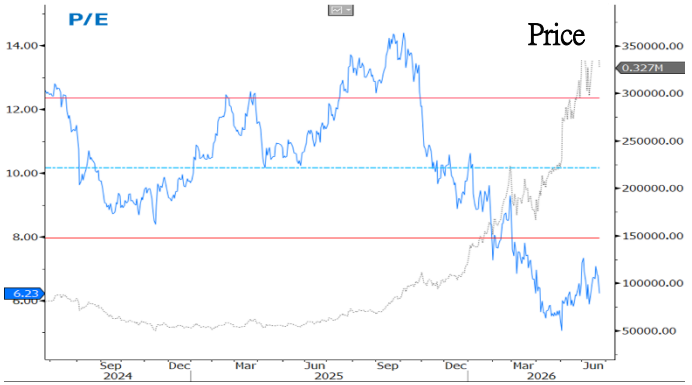

Storage will drive revenue growth

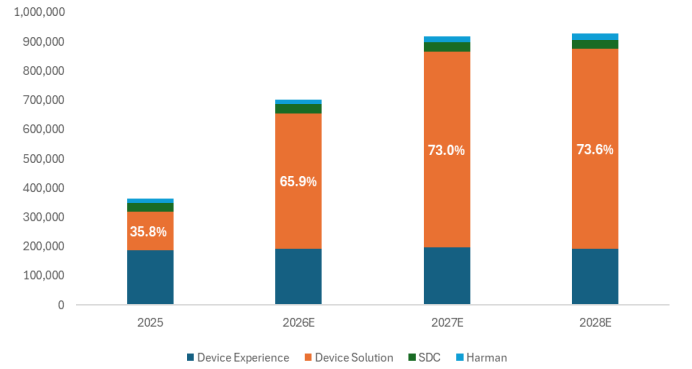

Volatility

Hist. Vol Imp. Vol

EU Stocks: ASML (ASML)

1. Investment Thesis

As the sole leader in lithography systems, ASML enjoys exceptional capacity and technology advantages. Data indicate its 2027 capacity can support >50% YoY growth in leading-edge wafer output, far above market demand of 25–30%; on DRAM, the 3800F delivers up to 87% higher throughput vs the 3600D. Benefiting from scaling (die count from 1b to 1c up 30%), the memory revenue mix is set to reach 30–35% in 2026, while DRAM EUV layers at the 1d node will jump to eight, lifting lithography intensity to 15–20%. With manufacturing lead time reduced from 20 to 14 weeks, Low NA EUV shipments are expected to reach 90/105 units in 2027/28, underpinning strong growth momentum.

2. Valuation Summary

The current valuation is highly attractive, with the premium to US peers narrowing to about 6%, far below the 84% historical average. Based on earnings growth, the Street has raised the target price to EUR 1,599 per share. This implies only ~32x 2028E EPS, at the company's 5-year forward P/E average 33x, providing ample margin of safety. As 2027/28E EPS is revised up by 9–21% and 20–25% on robust lithography demand, re-rating momentum is well supported.

3. Share Price Catalysts

Key catalysts include: (1) EUV/DUV orders from leading logic and memory customers growing ahead of expectations over the next few quarters, confirming tight end-market capacity; (2) initial High NA lithography shipments and process validation progress, which are expected to deliver 20–40% cost savings and process simplification on critical layers; (3) a rebound in immersion DUV shipments—especially ex-China leading-edge capacity expansions—and China immersion shipments recovering to 85 units in 2027, back to the 2025 peak, serving as a near-term revenue driver.

4. Investment Risks

Key risks: (1) escalation in geopolitics and export controls—particularly restrictions on exporting lithography tools and after-sales services to key markets such as China—which would directly pressure revenue; (2) delayed technology adoption, if TSMC (TSM US / 2330 TT) and other major customers postpone deployment of costly High NA tools on pricing considerations, raising concerns about the durability of equipment demand; (3) high customer and industry concentration, where sharp swings in Capex by a few semiconductor manufacturing giants could materially impact results.

5. Company Profile

The company holds over 80% global share in lithography equipment and is the world's only EUV supplier. It is an indispensable cornerstone for scaling, advanced logic R&D, and next-generation DRAM at TSMC (TSM US / 2330 TT), Samsung Electronics (005930 KS), Intel (INTC), and SK Hynix (000660 KS).

Financials (EUR)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	28.3	32.7	39.1	48.0	53.7
Adj. EBIT (100 Mn)	9.0	11.3	14.3	18.9	22.0
Adj. EPS (EUR)	19.2	24.7	31.7	42.1	49.5
OCF (100 Mn)	11.2	12.7	10.7	17.1	20.2
FCF (100 Mn)	9.1	11.1	8.6	14.8	16.8
Margins and growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue growth	2.6%	15.6%	19.6%	22.8%	11.9%
Adj. EPS growth	-3.3%	28.4%	28.1%	32.9%	17.6%
Gross margin	51.3%	52.8%	52.5%	54.1%	55.3%
Adj. EBIT margin	31.9%	34.6%	36.5%	39.5%	41.0%
Net margin	26.8%	29.4%	30.9%	33.3%	34.3%
Net debt/equity	-43.2%	-45.2%	-54.7%	-85.8%	-127.1%
Return on capital	47.4%	50.5%	54.4%	57.8%	51.8%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	1.4%	1.7%	1.3%	2.3%	2.6%
Dividend yield (%)	0.1%	0.1%	0.5%	0.6%	0.7%
EV/EBITDA(x)	26.0	28.0	41.7	31.7	27.4
P/E (x)	35.3	37.3	52.4	39.4	33.5
P/B (x)	14.4	18.1	26.7	19.9	14.9
P/S (x)	9.4	11.0	16.5	13.4	12.0

Volatility

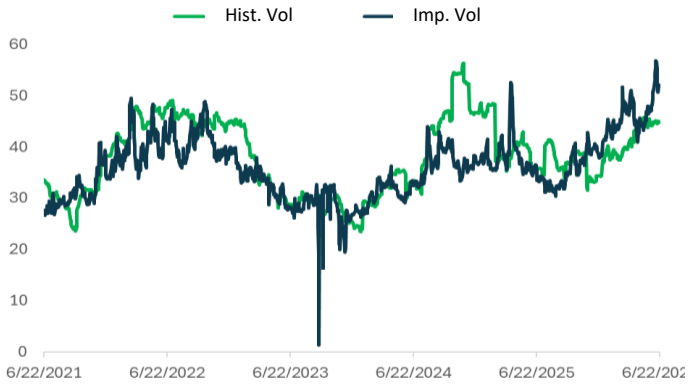

Valuation

Delivery by type

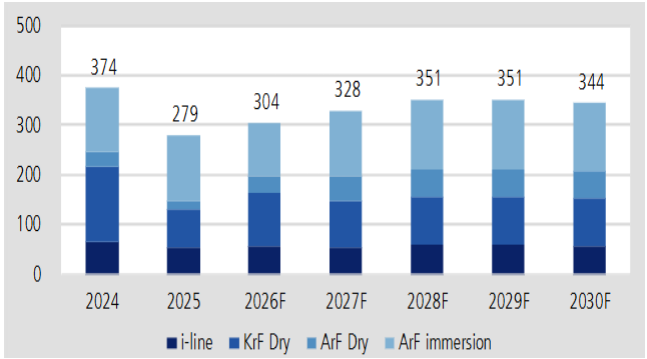

US equities: Alphabet (GOOGL.US)

1. Investment rationale (Investment Thesis)
The core ads franchise remains exceptionally solid, with Search and YouTube ad revenue accelerating sharply in 1Q. The Cloud business has strong margin-upgrade headroom, supported by core cloud service price hikes and royalties recognized from external sales/licensing of in-house TPU chips, driving Cloud OPM toward the 35%–42% levels of the industry leader and successfully pivoting to an asset-light, high-margin, diversified AI monetization model.

2. Valuation Summary
Target price raised to USD 433, based on 28x 2027E EPS. Consensus has largely priced in about USD 377 billion of AI Capex build over the next two years. As in-house chips and cloud services enter a high-ROI payback phase, the unique AI optionality and monetization potential of its massive consumer ecosystem (Search, Maps, YouTube) will drive a steady multiple re-rating.

3. Share Price Catalysts

Key catalysts include: (1) flagship Gemini 3.5 Pro launched in June, addressing prior gaps on long-horizon tasks and Harness Engineering benchmarks; (2) tests commencing for the 24/7 personal agent Gemini Spark and Antigravity 2.0 interactive endpoints showcased at the developer conference; (3) cloud price increases flowing through and royalties from external TPU sales/licensing recognized.

4. Investment Risks
Potential risks include: (1) a pricing shift from daily-usage caps to compute-based metering may trigger pushback and churn among individual and creator users; (2) on Agentic Coding and complex multi-step decision-making, Gemini currently trails closed-source competitors on some metrics; (3) a macro slowdown could suppress enterprise software spend, alongside rising global antitrust and regulatory scrutiny.

5. Company Profile
The company is the global leader in web search, digital advertising, and software ecosystems, with multiple flagship products exceeding 1 billion users (Search, YouTube, Maps). It also operates a top-three global cloud infrastructure platform and has full-stack development capabilities spanning in-house processors (TPU) and the multimodal Gemini models, pivoting toward the Agent Era to reshape digital productivity.

Financials (USD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	295.1	342.9	420.9	512.0	603.2
Adj. EBIT (100 Mn)	112.4	129.0	170.0	207.8	249.8
Adj. EPS (USD)	9.9	12.9	14.3	15.3	18.4
OCF (100 Mn)	125.3	164.7	210.6	259.9	324.2
FCF (100 Mn)	72.8	73.3	20.45	14.6	53.4
Margins and growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue growth	15.1%	16.2%	22.7%	21.7%	17.8%
Adj. EPS growth	70.2%	30.2%	11.1%	7.4%	20.0%
Gross margin	58.2%	59.7%	70.3%	69.6%	69.1%
Adj. EBIT margin	32.1%	32.0%	40.4%	40.6%	41.4%
Net margin	28.6%	32.8%	35.8%	32.0%	31.4%
Net debt/equity	-21.5%	-15.4%	-15.9%	-15.2%	-26.9%
Return on capital	32.9%	35.7%	55.5%	41.7%	57.4%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	1.6%	1.6%	0.5%	0.3%	1.2%
Dividend yield (%)	0.1%	0.1%	0.2%	0.2%	0.2%
EV/EBITDA(x)	17.0	24.5	19.3	15.3	12.3
P/E (x)	23.5	29.0	25.8	24.0	20.0
P/B (x)	7.1	9.2	7.8	6.0	4.6
P/S (x)	6.7	9.4	10.6	8.7	7.4

Volatility

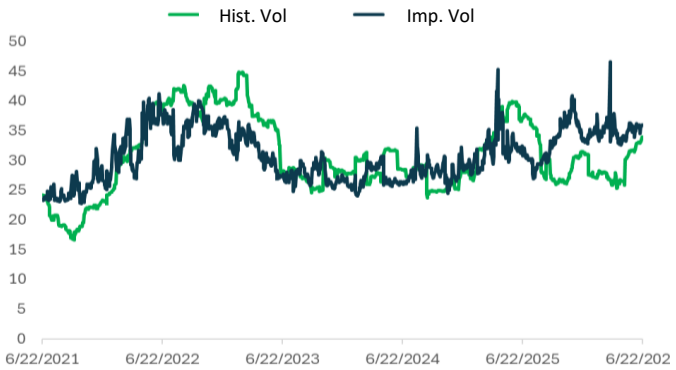

Valuation breakdown

CAPEX / Revenue

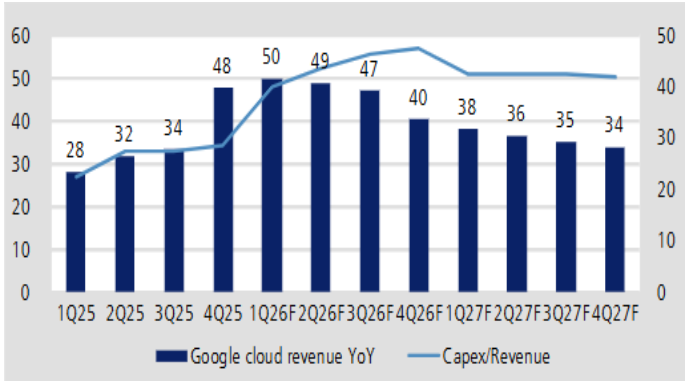

US stocks: Apple (AAPL)

1. Investment Rationale (Investment Thesis)

1Q revenue surged 17% to USD 111.2 billion, with iPhone ASP topping USD 1,000 for the first time in a calendar 1Q, underscoring strong pricing power. High-margin Services also reached USD 31.0 billion. Backed by a sticky ecosystem, long-term contracts to lock in memory capacity to mitigate raw-material price inflation, and a higher mix of high-end Pro models, the company's profitability continues to deliver robust growth despite industry headwinds.

2. Valuation Summary (Valuation Summary)

We raise our target price to USD 320, based on 32x P/E of FY2027E EPS, near the upper end of its historical range. With high-margin Services now contributing nearly half of profits, Apple deserves a premium valuation versus handset peers. As hardware shipments outpace the industry average and product mix improves, we expect earnings to grow steadily over the next two years.

3. Share Price Catalysts (Share Price Catalysts)

Key catalysts include: partnering with Google to launch a new-generation AI Siri with screen awareness; the budget iPhone 17e adopting an in-house C1X communications chip to save Qualcomm royalties, expected to save USD 1.15 billion by FY27; and rapid monetization of the Apple Creator Studio subscription service debuting in early 2026, initially likely to contribute USD 4.8 billion in annual revenue.

4. Investment Risks (Investment Risks)

Key downside risks include: a weaker-than-expected recovery in global smartphones; delays in the in-house chip program, or connectivity stability issues with the first-generation communications chip (e.g., intermittent WiFi dropouts) requiring time to optimize, affecting rollout timing across the lineup; additionally, continued memory price increases raising near- to medium-term hardware cost pressure, and potential global tariff policy changes that could further push up manufacturing supply-chain costs.

5. Company Profile (Company Profile)

The company is the world's most valuable personal electronics and consumer tech leader, specializing in designing, manufacturing, and selling smartphones, PCs, and wearables, and offering a rich suite of software and value-added services. Growth is driven by a powerful closed ecosystem of hardware and a high-margin (76.6%) services business. In recent years, the high-profit services segment's share of earnings has risen to nearly half, and the company is accelerating its AI strategy via on-device and private-cloud security architectures.

Financial data (USD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	391.0	416.2	478.1	520.6	557.0
Adj. EBIT (100 Mn)	123.2	133.1	155.1	167.2	181.3
Adj. EPS (USD)	6.8	7.5	8.8	9.7	10.7
FCF (100 Mn)	118.3	111.5	150.0	165.9	181.5
FCF (100 Mn)	108.8	98.8	139.8	151.1	167.6
Margins and growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue growth	2.0%	6.4%	14.9%	8.9%	7.0%
Adj. EPS growth	10.1%	10.5%	17.4%	10.2%	11.2%
Gross margin	46.2%	46.9%	48.2%	47.8%	48.3%
Adj. EBIT margin	31.5%	32.0%	32.4%	32.1%	32.6%
Net margin	24.0%	26.9%	26.8%	26.9%	27.2%
Net debt/equity	-66.0%	-27.2%	-95.8%	-128.8%	-177.4%
Return on capital	157.4%	171.4%	131.1%	108.4%	81.8%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	2.5%	2.3%	3.2%	3.5%	3.8%
Dividend yield (%)	0.1%	0.1%	0.4%	0.4%	0.4%
EV/EBITDA(x)	24.9	25.6	25.4	23.5	21.7
P/E (x)	37.5	34.2	34.0	30.9	27.8
P/B (x)	60.5	51.2	38.3	26.8	18.5
P/S (x)	8.9	9.2	9.2	8.4	7.9

Volatility

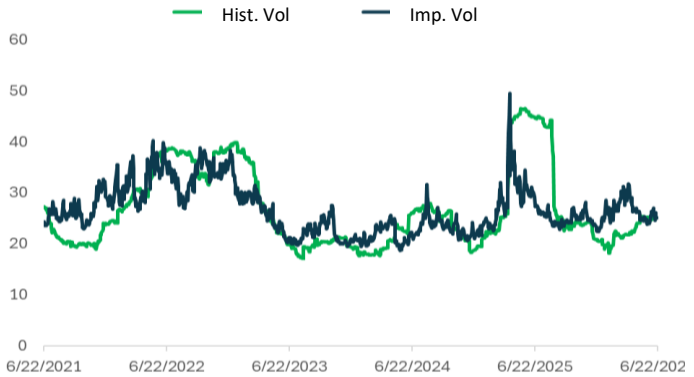

Valuation

Revenue Breakdown

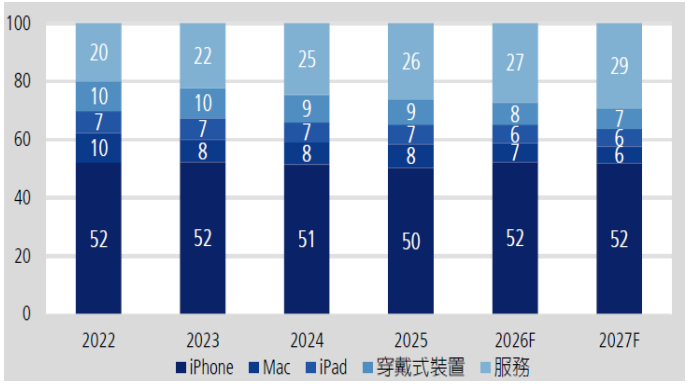

US equities: Microsoft (MSFT)

1. Investment Thesis

Azure, Microsoft's core AI growth engine, grew 39% in the third quarter; with new capacity coming online, growth is expected to accelerate further in 2H. In addition, paid M365 Copilot seats have surpassed 20 million, driving 15% growth in Commercial Cloud revenue. The company continues to expand proprietary MAI models and optimize Maia chips, effectively lowering operating costs and cementing full-stack AI leadership.

2. Valuation Summary

The market assigns Microsoft a target price of USD 561. The valuation premium is primarily underpinned by its leadership in GenAI and high-renewal-rate enterprise subscription revenue. While surging AI infrastructure investment is pressuring near-term FCF, discounting to FY2028E EPS of USD 22.8, the P/E remains attractive with very high long-term growth visibility.

3. Share Price Catalysts

Key catalysts include: (1) Azure growth re-accelerates in 2H on new capacity; (2) M365 Copilot users are expected to surpass 25 million in Q4; (3) the M365 E7 release launched in May integrates autonomous Agent AI, likely triggering upgrade waves in regulated industries (e.g., financials); (4) pricing is broadly shifting to usage/Token monetization.

4. Investment Risks

Key risks include outsized AI Capex (calendar 2026 guidance at USD 190 billion); if AI adoption lags expectations, margins will be dragged. Second, hardware inflation pressures cloud GM (expected to fall to 64%). In addition, competition is intense in the hyperscale cloud market, and the MPC segment faces component-cost and inventory overhang headwinds, constraining near-term recovery.

5. Company Profile

Microsoft (MSFT) was founded in 1975 and is the world's largest software leader. Starting with Windows and Office, it has in recent years successfully transformed into a cloud- and SaaS-driven giant. Azure has evolved into a top-tier global cloud infrastructure platform, and with AI full-stack capabilities spanning in-house Maia chips and Azure OpenAI to front-end Copilot applications, it delivers secure and compliant intelligent Agent solutions for enterprises. This positions Microsoft as the most critical and indispensable IT leader in enterprise automation and digital transformation in the GenAI era.

Financial Data (USD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	245.1	281.7	329.5	384.0	452.8
Adj. EBIT (100 Mn)	109.4	128.5	153.6	180.1	211.5
Adj. EPS (USD)	11.8	13.6	17.1	19.5	22.8
OCF (100 Mn)	118.5	136.2	174.6	213.8	259.4
FCF (100 Mn)	74.1	71.6	63.4	32.3	58.3
Margins and growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue growth	15.7%	14.9%	16.9%	16.5%	17.9%
Adj. EPS growth	21.9%	15.6%	25.5%	13.8%	17.2%
Gross margin	69.8%	68.8%	67.8%	66.6%	65.7%
Adj. EBIT margin	44.6%	45.6%	46.6%	46.9%	46.7%
Net margin	36.0%	36.1%	40.9%	42.5%	44.8%
Net debt/equity	8.3%	5.1%	-13.2%	-8.3%	-7.5%
Return on capital	37.1%	33.3%	35.0%	27.7%	25.7%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	2.6%	2.5%	2.2%	1.1%	2.1%
Dividend yield (%)	0.2%	0.2%	1.0%	1.1%	1.2%
EV/EBITDA(x)	24.7	22.1	14.3	11.8	9.7
P/E (x)	37.9	36.5	22.2	19.5	16.6
P/B (x)	12.4	10.8	6.3	4.9	3.8
P/S (x)	13.5	13.1	8.6	7.3	6.2

Volatility

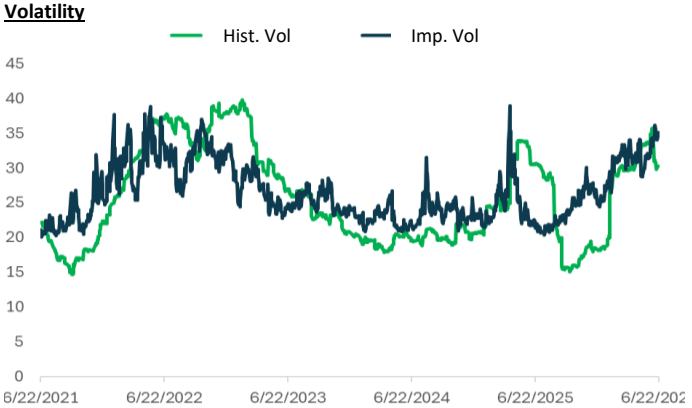

Valuation

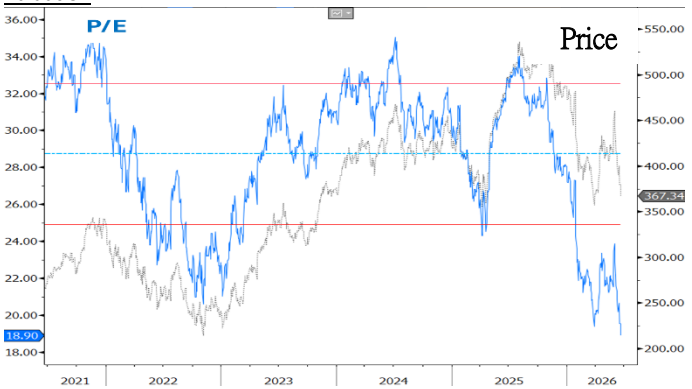

CAPEX / Revenue

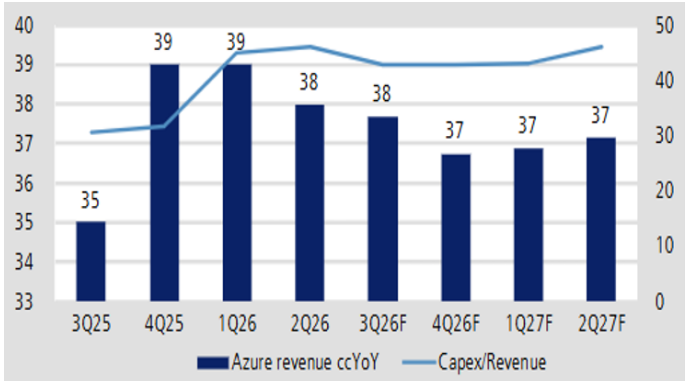

US: Amazon (AMZN)

1. Investment rationale (Investment Thesis)

Growth in core business AWS is accelerating; 1Q26 y/y +28%, to a 150B annualized revenue run-rate. Backlog surged ~50% q/q to 364B, driven by core workload migration and broad enterprise adoption of AI. Bedrock customer spend grew +170% q/q; in-house silicon is at a 20B annualized revenue run-rate and growing triple digits, with Trainium committed orders as high as 225B. E-commerce segment unit growth accelerated +300bps to +15%, a post-pandemic high, while essentials mix rose to ~40%, driving higher purchase frequency and basket size. The margin expansion trajectory is clear: North America retail operating margin reached 7.9%, 110bps above expectations—making it a scarce asset delivering high-quality earnings growth funded by internally generated FCF.

2. Valuation Summary

Consensus target prices cluster at USD 315, implying roughly 27x 2027E P/E—near the low end of the past decade's multiple range and highly defensive. With the current share price around USD 260, implied upside exceeds 20%. As AWS growth visibility improves, AI backlog unlocks, and retail efficiency drives rapid free cash flow (FCF) growth (expected to increase meaningfully), there is substantial room for multiple re-rating.

3. Share Price Catalysts

Key catalysts include: continued acceleration in AWS revenue growth, especially strong demand for Bedrock and GenAI platforms translating into realized revenue; Prime Day pulled forward to 2Q, expected to add about USD 7 billion incremental revenue; broad rollout of Amazon Supply Chain Services (ASCS), likely to attract more third-party sellers and lift returns on logistics assets; and further capacity ramp of in-house chips (Trainium, Graviton), bolstering confidence in its AI infrastructure competitiveness.

4. Investment Risks

Risks include: (1) Capex surging—remaining elevated at USD 200 billion in 2026 and reaching USD 291 billion in 2027; if AI monetization falls short, margins could be compressed; (2) macro slowdown weighing on consumer retail and advertising spend; (3) US/EU antitrust regulation (e.g., the EU Digital Markets Act) imposing operational constraints or high compliance cost pressure.

5. Company Profile

Amazon (AMZN) is a global leader in e-commerce and cloud, serving consumers, sellers, developers, and enterprises. Centered on Prime members, it has built an unmatched retail and last-mile logistics network, and via AWS provides cloud services and a GenAI stack with roughly 30% global market share.

Financials (USD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	638.0	716.9	824.9	933.4	1035.3
Adj. EBIT (100 Mn)	68.6	80.0	104.0	130.9	165.8
Adj. EPS (USD)	7.2	9.0	10.3	11.6	14.2
OCF (100 Mn)	115.9	139.5	183.7	236.6	296.1
FCF (100 Mn)	32.9	7.7	-15.8	-4.13	53.22
Margins and growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue growth	11.0%	12.4%	15.1%	13.2%	10.9%
Adj. EPS growth	146.6%	25.4%	14.9%	12.5%	22.6%
Gross margin	48.9%	50.3%	51.2%	52.6%	54.2%
Adj. EBIT margin	10.8%	11.2%	12.6%	14.0%	16.0%
Net margin	9.3%	10.8%	11.6%	11.7%	13.2%
Net debt/equity	16.3%	11.4%	10.7%	19.2%	30.8%
Return on capital	24.3%	22.3%	20.4%	18.0%	280.4%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	1.3%	0.3%	-0.6%	-0.2%	2.0%
Dividend yield (%)	0.0%	0.0%	0.0%	0.0%	0.0%
EV/EBITDA(x)	17.8	15.2	12.4	9.9	8.1
P/E (x)	39.7	32.2	23.7	21.1	17.2
P/B (x)	8.1	6.0	4.9	3.8	3.1
P/S (x)	3.6	3.4	3.2	2.8	2.5

Volatility

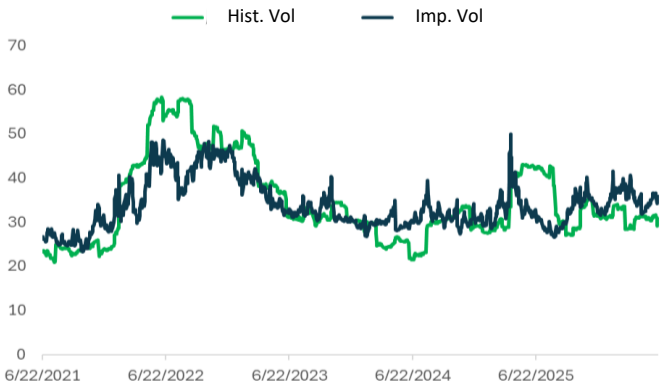

Valuation breakdown

CAPEX / Revenue

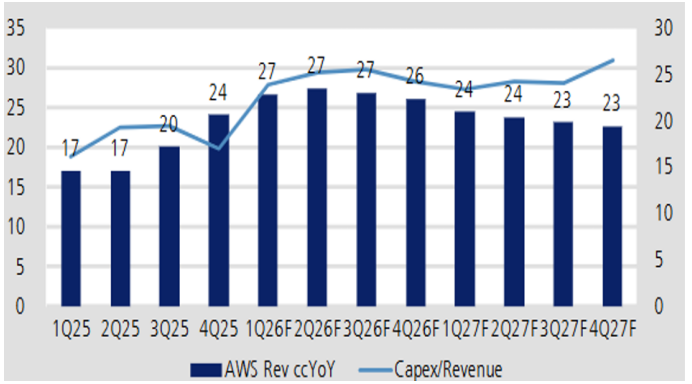

US: Meta Platforms (META)

1. Investment Rationale (Investment Thesis)
Meta's core ads and user engagement are showing strong AI-driven momentum. 1Q revenue reached USD 56.3 billion (+33% YoY, +29% FX-neutral); ad revenue was USD 55.0 billion. Algorithmic optimization, aided by "double-length interaction data sequences" and fast indexing, boosted stickiness: IG Reels time spent +10%, FB video views +8% QoQ. Lattice and GEM delivered >6% uplift in ad conversion; over 8 million advertisers have adopted GenAI tools; the value-optimization suite reached USD 20.0 billion in annual revenue. As commercial AI agents roll out globally (LatAm, Indonesia, etc.), messaging conversations have surged 10x to 10M+ per week, with monetization pathways increasingly clear.

2. Valuation Overview (Valuation Summary)
USD 819 target based on 21x 2027E EPS (EV/EBITDA ~12x). The current price implies ~17x 2027E P/E, and its massive global user network plus ad-tech premium offer solid valuation defensibility.

3. Share Price Catalysts (Share Price Catalysts)

Key catalysts include: global rollout of commercial AI agents and expansion to Instagram, accelerating AI monetization; WhatsApp launching native shopping and payments (no redirects), fostering chat-commerce habits and materially lifting messaging revenue; iterations and optimization of Muse Spark LLM technology—if they continue to enhance ad ROI, earnings expectations are likely to be revised up.

4. Investment Risks (Investment Risks)
Key risks include hefty AI spend with unclear payoff. The 2026 budget has been raised to USD 125–145 billion, pressuring FCF. LLM competition is intensifying and monetization beyond ads remains difficult. In addition, new EU privacy rules constrain targeting, and geopolitical conflicts weigh on advertisers' spending appetite, potentially pressuring near-term ad revenue growth.

5. Company Profile (Company Profile)
The company is a global social and communications giant; Facebook, Instagram, and WhatsApp have built an ecosystem covering billions of users worldwide. Its core business is driven by network effects, monetizing primarily via targeted ads. The company is pursuing a full AI transformation, deeply integrating large-model technologies and actively expanding new messaging monetization channels such as business AI agents.

Financials (USD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	164.5	201.0	253.1	301.7	353.5
Adj. EBIT (100 Mn)	69.4	83.3	90.6	105.1	123.5
Adj. EPS (USD)	30.2	31.4	40.2	38.5	46.5
OCF (bn)	91.3	115.8	137.3	166.8	209.2
FCF (bn)	54.1	46.1	2.38	2.76	27.66
Margins and growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue growth	21.9%	22.2%	25.9%	19.2%	17.2%
Adj. EPS growth	103.4%	3.9%	27.9%	-4.1%	20.6%
Gross margin	81.7%	82.0%	80.3%	79.1%	78.2%
Adj. EBIT margin	42.2%	41.4%	35.8%	34.8%	35.0%
Net margin	37.9%	30.1%	33.6%	29.8%	30.0%
Net debt/equity	-15.4%	1.6%	-1.6%	3.9%	1.8%
Return on capital	37.1%	30.2%	31.6%	25.0%	21.2%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	3.7%	3.1%	0.2%	0.2%	1.9%
Dividend yield (%)	0.1%	0.1%	0.4%	0.4%	0.4%
EV/EBITDA(x)	16.6	16.0	11.5	9.4	7.5
P/E (x)	24.5	28.1	14.4	15.0	12.4
P/B (x)	8.1	7.7	4.8	3.8	3.1
P/S (x)	9.0	8.3	5.8	4.9	4.1

Volatility

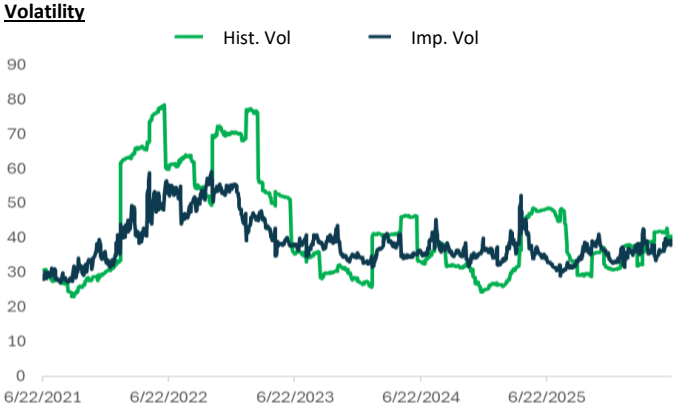

Valuation Breakdown

CAPEX / Revenue

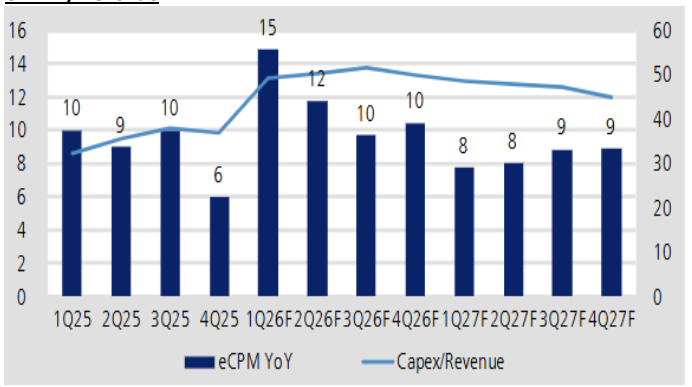

US Equities: Tesla (TSLA.US)

1. Investment Thesis

Despite Tesla facing volume headwinds in 1Q26 (deliveries down 14% q/q to 358,000 units), a one-off tariff refund drove GM to 21.1%, well ahead of expectations, with Non-GAAP EPS at USD 0.41, underscoring operating resilience. To cement its leadership in physical AI, the company raised 2026 Capex to USD 25 billion, including USD 3 billion for a self-built R&D fab in Texas to accelerate foundational compute for AI and the new Optimus lines. In addition, FSD subscriptions reached 1.28 million (attach rate 14%), and Netherlands approval should effectively advance EU rollout; we estimate the attach rate to reach 23% in 2027. Combined with year-end expansion of Robotaxi service coverage, high-margin software could reignite long-term growth.

2. Valuation Summary

Rating: Add. PT cut to USD 422. Valuation uses SOTP; factoring expanded AI investment turning near-term cash flow negative, we lower 2026–2028 EPS to USD 1.9 / 2.4 / 3.3, while the long-term physical AI leader premium remains attractive.

3. Share Price Catalysts

Catalysts include: (1) Robotaxi service expanding to 12 US states by year-end; (2) FSD gaining Netherlands approval to accelerate EU rollout; (3) FSD V15 upgrade by year-end enabling AI4 models to achieve unsupervised FSD; (4) Optimus V3 entering production by year-end with volume ramp in 2027.

4. Investment Risks

Key risks: (1) Ongoing EV price wars compress GM and profitability; (2) Capex raised to USD 25 billion in 2026 pressures FCF to turn negative; (3) Robotaxi/FSD rollout faces geopolitical and local regulatory approval delays.

5. Company Profile

Founded in 2003 and headquartered in Texas, Tesla is a global leader in EVs and physical AI. It began operating its first California plant in 2010 and has Gigafactories in Shanghai, Berlin, and Texas. Beyond its EV core, the ecosystem includes solar and energy storage (SolarCity), FSD, and the humanoid robot Optimus.

Financials (USD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	97.7	94.8	102.6	118.2	137.7
Adj. EBIT (100 Mn)	7.1	4.4	5.6	8.1	11.8
Adj. EPS (USD)	2.4	1.7	1.9	2.4	3.3
OCF (100 Mn)	14.9	14.7	15.3	17.9	22.5
FCF (100 Mn)	3.6	6.2	-9.6	-3.5	-0.6
Margins & Growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue growth	0.9%	-2.9%	8.2%	15.2%	16.5%
Adj. EPS growth	-43.7%	-31.4%	12.9%	29.9%	37.3%
Gross margin	17.9%	18.0%	19.6%	20.1%	21.2%
Adj. EBIT margin	7.2%	4.6%	5.4%	6.9%	8.6%
Net profit margin	7.3%	4.0%	7.1%	7.6%	8.8%
Net debt/equity	-31.1%	-35.4%	-25.6%	-21.7%	-14.5%
Return on capital	10.5%	4.9%	6.6%	7.9%	9.5%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	0.2%	0.4%	-0.6%	-0.2%	0.0%
Dividend yield (%)	0.0%	0.0%	0.0%	0.0%	0.0%
EV/EBITDA(x)	90.9	129.2	99.8	78.1	60.3
P/E (x)	198.0	416.4	213.7	164.5	119.8
P/B (x)	17.8	20.5	15.3	14.1	12.8
P/S (x)	13.2	15.3	14.7	12.7	10.9

Volatility

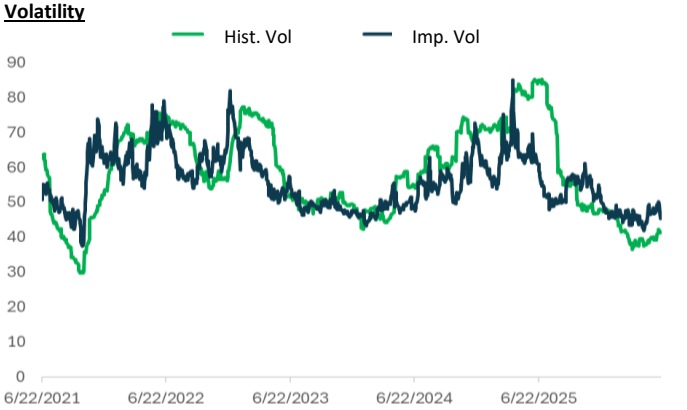

Valuation

CAPEX

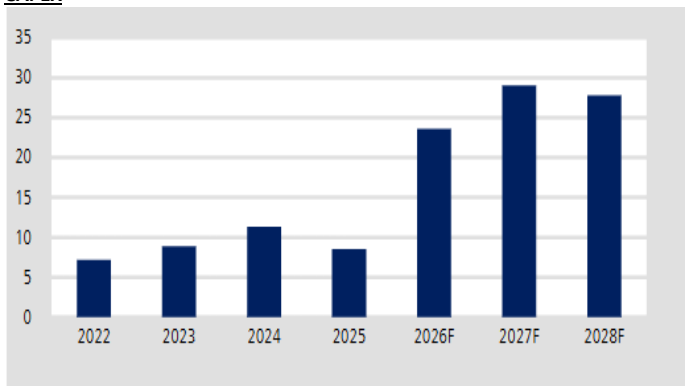

US Equities: AMD (AMD)

1. Investment Rationale (Investment Thesis)
Server CPU and AI businesses are at a structural inflection point. The Data Center CPU market is forecast to double to USD 120 billion by 2030, with CAGR >35%. Powered by the Turin chip, quarterly server CPU revenue grew >50% YoY, market share climbed to 44%, driving full-year segment growth of 80%. On AI GPUs, the first 1GW customized MI450 deployment under the 6GW agreement with Meta is expected to contribute up to USD 15 billion. We estimate AI sales to reach USD 33 billion in 2027, driving 2028 EPS to exceed USD 20.

2. Valuation Overview (Valuation Summary)
Valuation primarily uses SOTP. Data center GPU and CPU are valued at 35x and 30x 2028 P/E, respectively. Including client and other businesses plus about USD 35 per-share net cash, we raise the latest TP range to USD 498.

3. Share Price Catalysts
Top catalyst: 2H26 volume shipments of custom MI450 GPU to Meta (META) and large-model customers. Next, a surge in Agentic AI inference demand drives Turin and Venice server CPU deployments above expectations. Disclosure of a GW-scale partnership with a third hyperscaler (e.g., Microsoft (MSFT)) would further allay market concerns and trigger a re-rating.

4. Investment Risks
Risks: (1) Consumer markets (PC and gaming) face memory price inflation, with gaming revenue expected to decline >20% in 2H26; (2) mass production of custom rack systems for hyperscalers faces supply-chain and technical execution challenges; (3) initial GPU shipments are margin-dilutive, and together with the warrants' contra-revenue impact could cap near-term operating margin expansion.

5. Company Profile
The company is a leading fabless semiconductor giant, designing and selling x86 CPU for PCs and servers, GPU, and embedded and semi-custom chips. It has approximately 26,000 employees worldwide. Its product lines include Ryzen, EPYC, and Instinct AI accelerators. Leveraging architectural leadership and foundry partnerships, it has become a key leader in the global data center market.

Financials (USD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	25.8	34.6	49.8	76.5	101.1
Adj. EBIT (100 Mn)	6.1	7.7	13.5	24.6	34.8
Adj. EPS (USD)	3.3	4.2	7.2	13.1	17.6
OCF (100 Mn)	3.0	7.7	9.9	18.1	24.3
FCF (100 Mn)	2.4	6.7	8.6	16.7	24.3
Margins and growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue growth	13.7%	34.3%	43.7%	53.8%	32.1%
Adj. EPS growth	24.9%	26.0%	73.8%	80.9%	34.2%
Gross margin	49.4%	49.5%	55.5%	55.6%	55.8%
Adj. EBIT margin	23.7%	22.2%	27.1%	32.2%	34.5%
Net margin	21.0%	19.7%	24.5%	28.7%	29.1%
Net debt/equity	-5.1%	-10.6%	-19.4%	-42.8%	-82.0%
Return on capital	2.9%	7.2%	12.7%	18.2%	20.7%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	0.3%	0.8%	1.0%	1.9%	2.8%
Dividend yield (%)	0.0%	0.0%	0.0%	0.0%	0.0%
EV/EBITDA(x)	36.8	48.6	63.0	35.8	26.3
P/E (x)	125.2	81.1	74.1	41.0	30.5
P/B (x)	3.5	5.6	12.3	10.4	7.9
P/S (x)	7.9	10.1	17.6	11.4	8.7

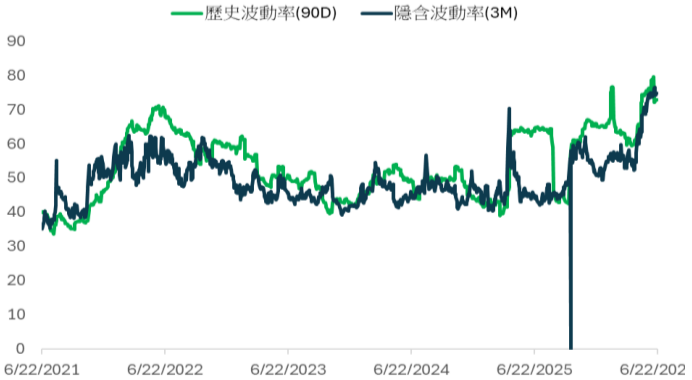

Valuation breakdown

Revenue Breakdown

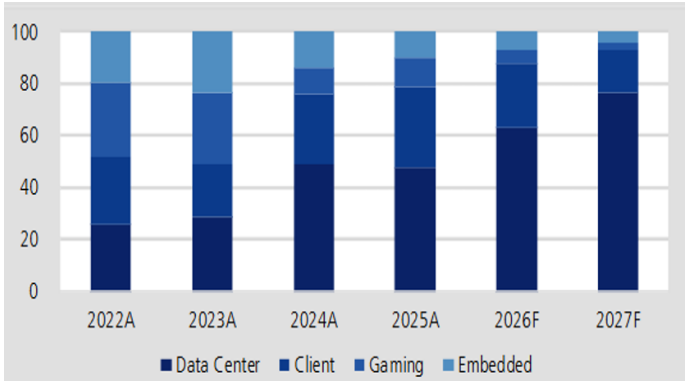

Volatility

Hist. Vol Imp. Vol

Japan equities: Renesas (6723 JP)

1. Investment Thesis (Investment Thesis)
Benefiting from the boom in AI inference, memory interface sales are expected to surge from JPY 110 billion in FY2025 to JPY 220 billion in FY2028. Data indicate global AI CPU shipments will grow 56% to 36 million units over 2025–2027. Together with industry price hikes (PMIC +10% expected), JPY depreciation, and JPY 400 billion after-tax gains from the sale of the clock business, operating profit is projected to rise to JPY 572 billion and JPY 666 billion in FY2026 and FY2027, respectively, underscoring a robust outlook.

2. Valuation Summary (Valuation Summary)
Valuation is primarily based on EV/EBITDA, a target price of JPY 4,796.

3. Share Price Catalysts (Share Price Catalysts)
Core catalysts include the upcoming Capital Markets Day on June 25, when detailed plans for AI servers and the newly acquired Altium business are expected. Additionally, outperformance in non-auto segments, expansion of the AI data center business, upgrades to DDR5 and DDR6 memory interfaces, and six consecutive quarters of QoQ revenue growth should bolster investor confidence and drive valuation re-rating.

4. Investment Risks (Investment Risks)
Key risks include: (1) weak global auto sales with Tier 1 customers maintaining cautious procurement; (2) rising energy prices accelerating the shift to EVs, potentially pressuring the company's market share; (3) increases in fixed costs such as R&D compressing margins; (4) macroeconomic slowdown and sharp JPY appreciation posing FX risks.

5. Company Profile (Company Profile)
The company is a global leader in specialty semiconductors, holding a leading position in microcontrollers (MCU) with roughly 20% global share. Its business is mainly split between Automotive and Industrial/Infrastructure/IoT. Revenue is primarily driven by microcontrollers (43%) and analog (35%). It is actively expanding into high-growth AI data center businesses such as PMIC and memory interfaces.

Financials (JPY)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (Bn)	1348.5	1321.2	1538.3	1679.5	1826.1
Adj. EBIT (Bn)	397.9	386.9	468.4	540.7	596.9
Adj. EPS (JPY)	201.5	182.3	248.2	259.8	293.2
OCF (Bn)	340.5	452.9	474.8	544.5	590.7
FCF (Bn)	-943.6	328.2	345.1	379.6	418.1
Margins and Growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue Growth	-8.2%	-2.0%	16.4%	9.2%	8.7%
Adj. EPS Growth	-17.3%	-2.8%	21.1%	15.4%	10.4%
Gross Margin	55.6%	57.1%	58.2%	59.0%	59.0%
Adj. EBIT Margin	29.5%	29.3%	30.4%	32.2%	32.7%
Net Profit Margin	26.7%	24.9%	28.1%	29.1%	30.3%
Net Debt/Equity	47.0%	38.0%	23.4%	7.0%	-13.1%
Return on Capital	9.7%	-2.1%	14.5%	11.9%	12.4%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF Yield (%)	2.4%	4.1%	3.9%	4.3%	4.7%
Dividend Yield (%)	-	-	0.6%	0.7%	0.8%
EV/EBITDA(x)	11.1	11.5	18.6	17.1	15.9
P/E (x)	16.9	-	19.2	18.3	16.2
P/B (x)	1.4	1.6	3.1	2.8	2.5
P/S (x)	2.7	2.9	5.8	5.3	4.9

Volatility

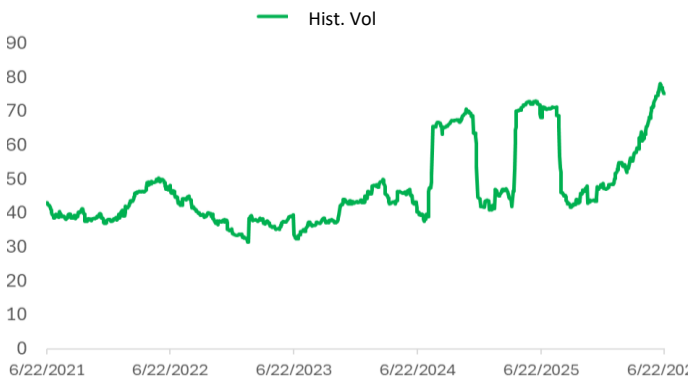

Valuation Breakdown

Revenue Breakdown

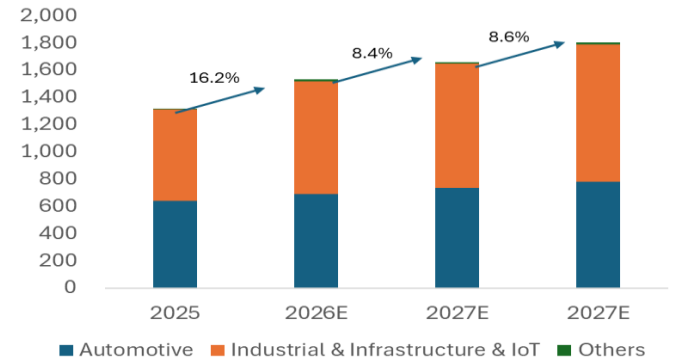

US stock: SanDisk (SNDK)

1. Investment Rationale (Investment Thesis)
Benefiting from strong AI tailwinds, the company's eSSD business surged over 200% QoQ in 1Q, lifting revenue mix to 25%. The New Business Model (NBM) is delivering: five LTAs signed; the first three have contract value (RPO) of up to USD 42 billion, locking in over one-third of FY2027 bit shipments, with USD 11 billion in financial guarantees that ensure floor GM above 80%. With industry supply/demand tightness, we expect 2026 NAND and eSSD ASPs to jump 186% and 265%, respectively, underpinning highly visible earnings growth.

2. Valuation Summary (Valuation Summary)
Given the high visibility provided by the company's unique long-term contract mechanism, the market has raised its target price from USD 1,845, implying 9–10x FY2027E P/E. Versus peers (e.g., JV partner at ~7x), the company deserves a premium. In addition, replacement Capex embedded in its JV fabs but not reflected on the financial statements amounts to USD 15–20 billion, implying substantial hidden value.

3. Share Price Catalysts (Share Price Catalysts)
Near-term key catalysts include: (1) advanced BiCS8-architecture NAND flash entering accelerated ramp and customer qualification; (2) the Board has approved up to USD 6 billion in open-market share repurchases (~3% of market cap), with each 1% reduction in share count expected to add roughly USD 2 to EPS; (3) a sharp rise in eSSD pricing driven by an H2 acceleration in AI infrastructure build-out.

4. Investment Risks (Investment Risks)
Key risks include: (1) a sharper-than-expected global macro slowdown that curbs cloud customers' data center Capex; (2) price wars initiated by other major players (especially China-based vendors aggressively chasing share), or a rapid near-term release of underutilized capacity, which could flip supply-demand and pressure product margins.

5. Company Profile (Company Profile)
Founded in 1988 and headquartered in California, USA, the company was acquired by Western Digital (WDC) in 2016 and completed a spin-off and standalone listing in February 2025. As a global NAND flash leader, it has extended its manufacturing JV with Kioxia to 2034, holding a 49% stake. It focuses on advanced NAND storage solutions for PCs, mobile, and—most critically—cloud data centers.

Financials (USD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	6.7	7.4	19.8	45.6	49.5
Adj. EBIT (100 Mn)	-0.5	0.7	11.8	35.9	38.2
Adj. EPS (USD)	-	3.0	65.6	191.3	204.5
OCF (100 Mn)	-	0.1	8.8	29.5	32.1
FCF (100 Mn)	-0.3	-0.2	8.4	27.9	31.4
Margins and Growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue Growth	9.5%	10.4%	169.0%	130.5%	8.7%
Adj. EPS Growth	-	-	1618.3%	203.1%	6.4%
Gross Margin	23.7%	30.3%	69.7%	82.6%	81.4%
Adj. EBIT Margin	-7.0%	9.4%	59.8%	78.7%	77.1%
Net Margin	-	6.0%	51.6%	66.1%	63.0%
Net Debt/Equity	-	6.4%	-82.3%	-367.6%	-561.4%
Return on Capital	0.0%	-16.2%	53.7%	61.5%	40.6%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF Yield (%)	-0.1%	0.0%	2.5%	8.3%	9.3%
Dividend Yield (%)	-	-	0.0%	0.0%	0.0%
EV/EBITDA(x)	-	11.2	27.6	9.7	9.9
P/E (x)	-	-	34.6	11.9	11.1
P/B (x)	-	0.7	18.2	7.1	4.3
P/S (x)	-	0.9	17.0	7.4	6.8

Volatility

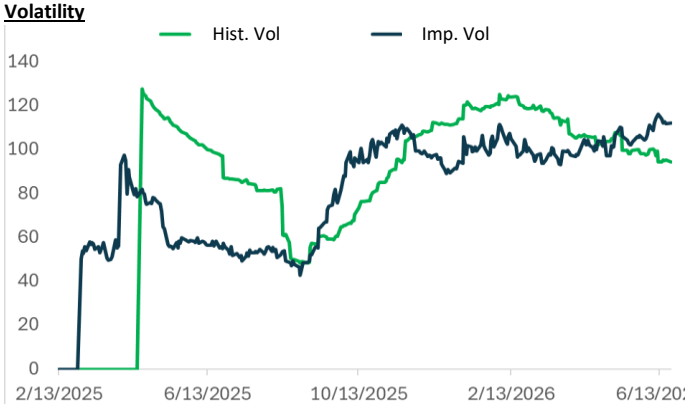

Valuation

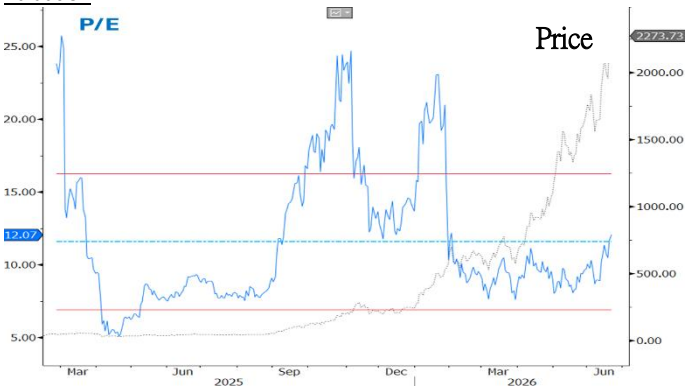

Productivity

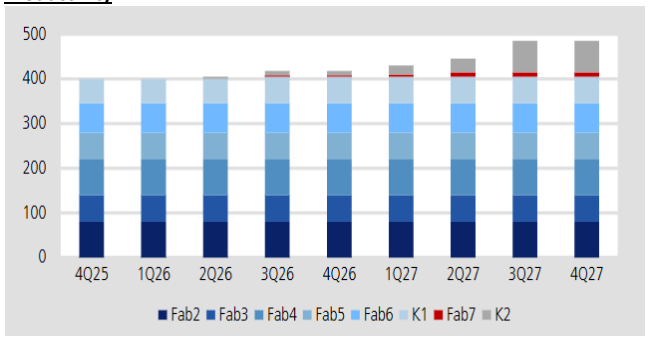

US Stocks: Marvell (MRVL)

1. Investment Thesis

FY2027 Q1 revenue reached USD 2.42 billion (+28% YoY), with data center revenue at USD 1.83 billion (+27% YoY). Leveraging leadership in optical interconnects and 51.2T switches, FY2027 data center revenue growth is raised to 50%, with optical interconnect up 70%, and a maintained 60–65% share in the 1.6T DSP market. As custom ASIC ramps in 2H, FY2028 (CY27) total revenue is estimated at USD 16.5 billion, with data center revenue further accelerating +55%; the company is moving toward its FY2029 goal of USD 10 billion in custom ASIC.

2. Valuation Summary

The market primarily applies a P/E approach, assigning 37x to CY27E EPS (USD 6.9), implying a target price of USD 256. This is a premium to AI semiconductor peers (26–27x), reflecting its unique leadership in high-speed optical interconnects (DSP) and custom AI chips. With EPS expected to approach USD 8.0 in 2028, the longer-dated valuation looks more attractive, and over the long term EPS could exceed USD 10 exiting CY28.

3. Share Price Catalysts (Share Price Catalysts)

Near-term catalysts include Amazon (AMZN)'s next-gen Trainium 3 custom chip shipping in the July quarter and ramping in 2H, and Microsoft (MSFT)'s Maia 3nm chip expected to enter high-volume production in early next year. In addition, an expanding customer base for 51.2T switching chips, accelerated penetration of 1.6T DSP, and high-speed eSSD storage controllers surpassing USD 1 billion in revenue in CY26 will drive upward estimate revisions. The company has already begun next-gen 2nm (e.g., Trainium 4) product designs and is actively pursuing companion XPU chip orders from hyperscale cloud customers, providing sustained share-price catalysts.

4. Investment Risks (Investment Risks)

Key risks include: (1) if cloud service providers slow Capex or AI data center builds, demand for optical interconnect and custom silicon will be directly suppressed; (2) intense competition in custom ASIC and optical communications; (3) advanced-node custom chips largely adopt HBM, and constrained advanced substrate supply could temporarily impact shipments.

5. Company Profile (Company Profile)

The company is a global leading fabless semiconductor solutions provider; since acquiring Inphi in 2021 it has cemented its leadership in high-speed optical communications. It focuses on four key markets—AI/cloud data centers, telecom infrastructure, enterprise networking, and automotive/industrial—and has become a core hardware pillar for global AI and high-performance computing.

Financial Data (USD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	5.8	8.2	11.5	16.8	23.4
Adj. EBIT (100 Mn)	1.7	2.9	4.2	6.5	9.6
Adj. EPS (USD)	1.6	2.8	4.0	6.9	9.2
OCF (100 Mn)	1.7	1.8	3.4	5.3	8.0
FCF (100 Mn)	1.4	1.4	2.76	4.55	6.78
Margins and Growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue Growth	4.7%	42.1%	40.6%	45.5%	39.4%
Adj. EPS Growth	-25.9%	80.9%	42.4%	71.1%	33.1%
Gross Margin	61.0%	59.5%	58.6%	57.5%	56.2%
Adj. EBIT Margin	28.8%	35.3%	36.4%	38.5%	41.1%
Net Margin	23.9%	30.1%	31.9%	34.2%	36.3%
Net Debt/Equity	25.3%	15.0%	-3.5%	-25.9%	-58.2%
Return on Capital	-6.3%	19.3%	12.0%	12.8%	26.6%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	0.5%	0.5%	1.0%	1.7%	2.5%
Dividend yield (%)	0.0%	0.0%	0.1%	0.1%	0.1%
EV/EBITDA(x)	71.8	81.8	58.8	40.6	27.3
P/E (x)	-	25.7	76.8	44.9	33.7
P/B (x)	7.3	4.7	15.3	13.0	9.7
P/S (x)	16.9	8.3	23.6	16.2	11.6

Volatility

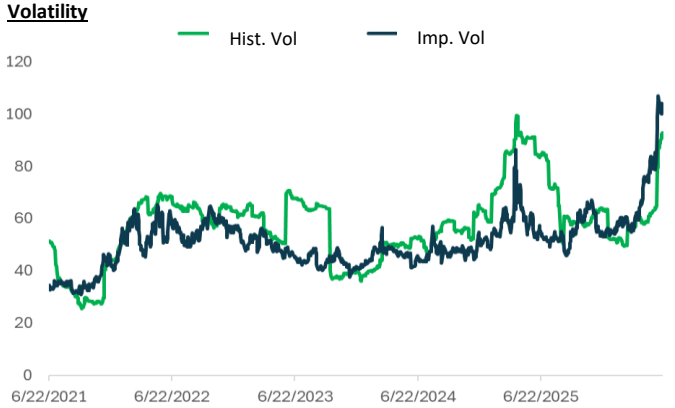

Valuation

Data Center revenue breakdown

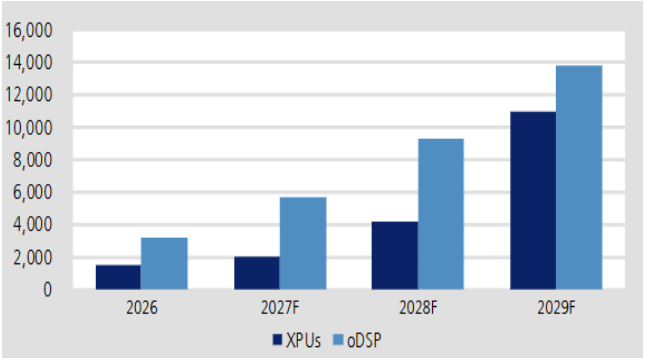

US stock: Coherent (COHR)

1. Investment Thesis

Benefiting from AI high-speed module demand, data center revenue grew 37% YoY in F3Q26, with communications up 60%. Capacity expansion is ahead of plan; critical indium phosphide (InP) capacity is expected to double by December 2026, one quarter early. The company has signed take-or-pay agreements with customers through 2028, providing high order visibility and low double-ordering risk. In addition, a Chinese competitor being placed on the US Department of Defense list is prompting US customers to de-risk, accelerating share shifts to the leading vertically integrated domestic player and underpinning growth with strong data center and geopolitical moats.

2. Valuation Summary

After the sector pullback, valuation looks more compelling; the stock trades at ~25x 2028E P/E. Against 40%+ earnings growth, PEG is <1x. The market primarily values the name on P/E, with a target price range of USD 385, implying ~32x FY2028 EPS. As AI-driven profit release gathers pace, valuation multiples have room to expand.

3. Share Price Catalysts

Key catalysts include 1.6T optical module capacity doubling over the next year with accelerated shipments; recent new orders from non-core large customers on CPO/NPO pluggable solutions; and a margin inflection with GM surpassing 40% by June 2026 on schedule, driven by yield improvements and pricing optimization. Order shifts materializing from geopolitical supply-chain reshuffling should also provide sustained momentum.

4. Investment Risks

Risks to watch include commoditization-led price pressure as optical technologies mature; potential for system OEMs to pursue vertical in-house development and reduce external sourcing; and integration risks from prior large M&A. Intensifying competition in AI and silicon carbide (SiC), and a macro downturn that weakens 5G and consumer demand, could also weigh on medium- to long-term GM and earnings growth.

5. Company Profile

Headquartered in Pennsylvania, USA, the company is a global leading vertically integrated manufacturer of optoelectronic devices and advanced materials. Its products are widely used in data center communications, industrial materials processing, aerospace & defense, and semiconductor equipment. With an end-to-end value chain from core wafers to optical modules, it provides optical modules, lasers, and optoelectronic solutions to major global data centers and OEM customers.

Financial data (USD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	4.7	5.8	7.1	9.6	12.7
Adj. EBIT (100 Mn)	0.1	0.3	1.4	2.2	3.1
Adj. EPS (USD)	1.7	3.5	5.5	8.3	12.0
OCF (100 Mn)	0.5	0.6	1.7	10.2	16.9
FCF (100 Mn)	0.2	0.2	-0.3	0.4	1.3
Margins and Growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue Growth	-8.8%	23.4%	21.5%	36.2%	32.4%
Adj. EPS Growth	-	111.4%	54.5%	52.6%	43.8%
Gross Margin	36.0%	37.9%	39.4%	40.5%	41.5%
Adj. EBIT Margin	2.0%	17.9%	20.4%	22.7%	24.5%
Net Margin	15.1%	0.8%	15.1%	17.7%	18.9%
Net Debt/Equity	42.5%	35.2%	10.9%	8.4%	-3.1%
Return on Capital	-4.7%	-1.5%	11.0%	15.2%	18.8%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF Yield(%)	0.3%	0.3%	-0.4%	0.6%	1.8%
Dividend Yield(%)	0.0%	0.0%	0.0%	0.0%	0.0%
EV/EBITDA(x)	18.9	16.6	45.0	30.9	21.4
P/E (x)	-	-	71.4	46.8	32.5
P/B (x)	2.1	2.5	7.2	6.6	5.7
P/S (x)	2.3	2.4	10.8	7.9	6.0

Volatility

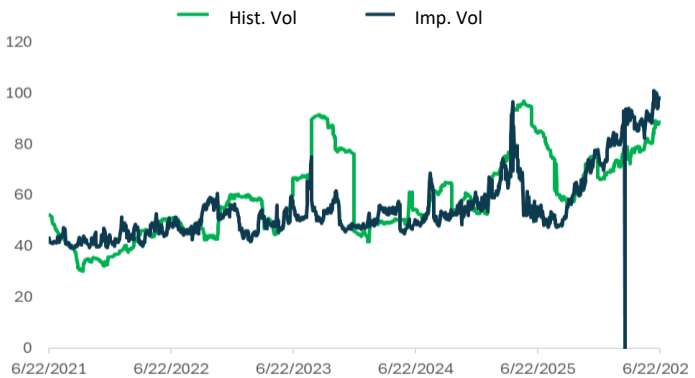

Valuation Breakdown

Revenue breakdown

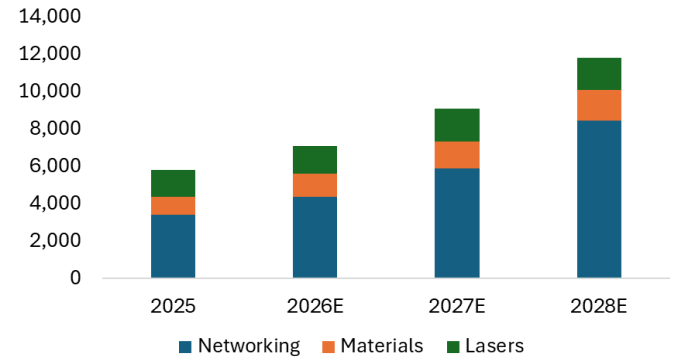

US equities: Lumentum (LITE)

1. Investment Thesis

F3Q26 revenue reached USD 808 million (+90% YoY), with Systems up 121%. On mix optimization and pricing, Non-GAAP GM expanded 540 bps QoQ to 47.9%, and EPS was USD 2.37. Strong next-quarter guide calls for revenue of USD 0.98–1.01 billion (+USD 180 million QoQ) and operating margin rising to 35–36%. The Cloud Light acquisition deepens Datacom exposure; at a key customer, CPO switch attach now exceeds 50%. As 100G/200G EML, OCS, and CPO ramp, quarterly revenue run-rate could reach USD 2.0 billion as early as Q3 2027, underscoring powerful growth momentum.

2. Valuation Summary

Although the share price has pulled back ~20% from its peak on near-term CPO adoption concerns, the current valuation is compelling. Implies only ~25x average P/E on CY28E EPS, with PEG <1x (EPS CAGR >40%), well below the historical premium. The Street's target prices are USD 1,126, based on CY28E EPS (~USD 30) with a 38x P/E, discounted back to end-2026.

3. Share Price Catalysts

Near-term upside catalysts include: (1) price hikes on supply-constrained Telecom and DCI optical components, directly boosting revenue and margins; (2) ramp of new external laser source (ELS) orders from non-mainstream AI chip customers; (3) accelerated Cloud Light module vertical integration to ease third-party CW laser supply bottlenecks. In addition, the U.S. placing major Chinese optical module rivals on restricted lists will prompt North American hyperscale data center customers to de-risk by accelerating order shifts to the company and other domestic, vertically integrated suppliers.

4. Investment Risks

Core risks include: (1) supply-chain bottlenecks from third-party CW laser constraints that could delay transceiver and chip deliveries; (2) customer concentration risk from high revenue exposure to key accounts (e.g., Google); (3) downstream optical system OEMs pursuing in-house vertical integration, weakening the company's pricing power.

5. Company Profile

The company is a global leader in optics and photonics, serving AI/ML, cloud, telecom, and industrial markets. It operates two segments: "Cloud and Networking," providing high-speed transceivers, OCS, and CPO components that power AI data center optical interconnects; and "Industrial Technology," focused on consumer electronics (e.g., 3D sensing), automotive, and industrial high-power lasers. Through M&A integration, the company has built strong, vertically integrated R&D and manufacturing advantages.

Financials (USD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	1.4	1.6	3.0	5.6	8.6
Adj. EBIT (100 Mn)	-0.4	-0.2	0.9	2.2	3.6
Adj. EPS (USD)	1.0	2.1	8.2	18.1	29.1
OCF (100 Mn)	0.0	0.1	0.5	1.5	2.7
FCF (100 Mn)	-0.1	-0.1	0.1	1.0	2.12
Margins and Growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue growth	-23.1%	21.0%	82.2%	85.9%	55.2%
Adj. EPS growth	-152.3%	104.0%	295.9%	121.6%	61.3%
Gross margin	33.0%	34.7%	45.3%	49.7%	51.6%
Adj. EBIT margin	2.8%	9.7%	29.1%	38.9%	41.5%
Net profit margin	-31.9%	-10.9%	29.1%	38.9%	41.5%
Net debt/equity	174.7%	152.6%	74.0%	-22.3%	-190.4%
Return on capital	-47.3%	2.5%	-	-	-
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	-0.2%	-0.2%	0.2%	1.6%	3.2%
Dividend yield (%)	0.0%	0.0%	0.0%	0.0%	0.0%
EV/EBITDA(x)	146.5	81.5	67.7	27.9	16.6
P/E (x)	-	256.1	104.2	47.0	29.2
P/B (x)	3.5	5.8	18.6	11.8	7.4
P/S (x)	2.5	4.0	22.1	11.9	7.6

Volatility

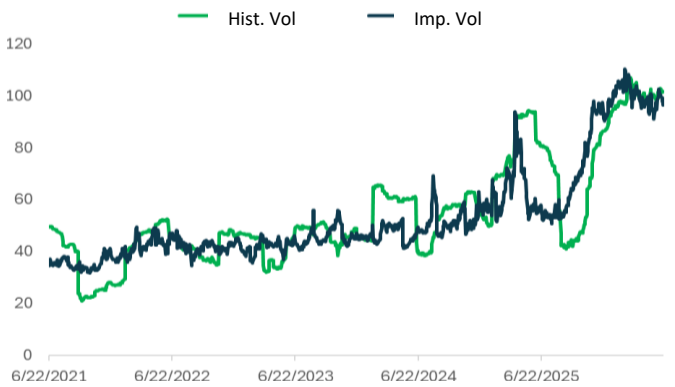

Valuation

Revenue breakdown

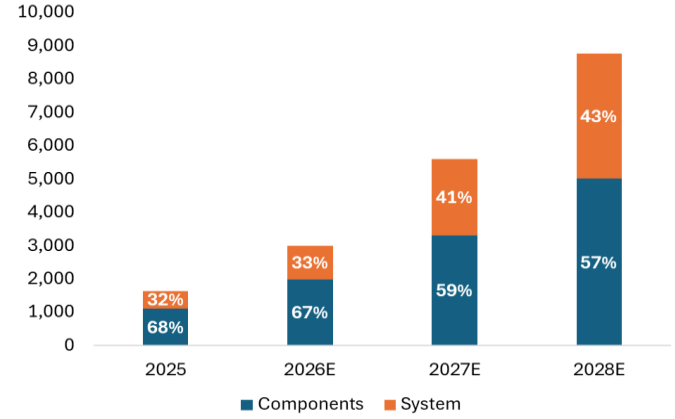

US Stock: GE Vernova (GEV)

1. Investment Thesis (Investment Thesis)
In 1Q, total orders reached USD 18.28 billion (book-to-bill 2.0x), driving backlog to USD 153 billion and pulling forward the USD 200 billion backlog target to 2027. Electrification orders were USD 7.1 billion; data center demand contributed USD 2.4 billion, exceeding last year's full-year level. Power segment orders reached USD 10 billion, with capacity booked out through 2028. With gas equipment pricing expected to be up 10% to 20% sequentially by 1H26 and data centers accounting for as much as 20% of power contracts, strong pricing power and AI-driven power demand provide high visibility and gross margin expansion runway.

2. Valuation Summary (Valuation Summary)
Valuation uses SOTP, with a fair target price range of USD 1,218; the current share price implies 15.5x 2029E EBITDA. The company commands a valuation premium given its leadership in the high-growth North American 60Hz market, an installed base twice that of peers, and rising pricing on high-margin service agreements. As core business margins rise, there is valuation re-rating potential ahead.

3. Share Price Catalysts (Share Price Catalysts)

Key catalysts include: solid-state transformer (SST) to enter testing in September, with commercial orders likely to commence in 1H27; the wind-related Section 232 tariff policy announcement will remove pricing uncertainty and accelerate ordering; in coming quarters, we expect sizable upfront deposits from nuclear customers, and the release of small modular reactor (SMR) orders in the US, Canada, and Sweden, which should significantly lift market sentiment.

4. Investment Risks
Key risks include: EPC contractor capacity bottlenecks could delay project delivery; if the turnaround and margin expansion in wind and electrification proceed slower than expected, overall results could be dragged; competition remains intense from peers such as Siemens Energy (ENR GR) and Vestas (VWS DC). In addition, changes in government energy policies, tariff increases beyond absorptive capacity, and any slowdown in the growth of data center AI power demand could pressure order growth.

5. Company Profile
The company is a global leader in the power sector, focused on power generation, transmission, conversion, and storage, with core businesses spanning generation, wind, and electrification. It benefits from a unique "Generation + Grid" integrated bundled-selling advantage, offering strong competitiveness amid electrification. With a global installed base roughly 2x larger than key peers, its high-margin long-term service agreements (LTSA) deliver highly visible, multi-decade stable cash flows.

Financials (USD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	34.9	38.1	45.4	51.9	59.2
Adj. EBIT (100 Mn)	0.5	1.4	4.6	7.6	10.2
Adj. EPS (USD)	5.6	17.7	25.2	23.8	32.4
OCF (100 Mn)	2.6	5.0	53.8	56.0	66.5
FCF (100 Mn)	1.7	3.7	6.5	7.1	9.17
Margins and Growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue Growth	5.1%	9.0%	19.4%	14.3%	14.0%
Adj. EPS Growth	-	217.0%	42.3%	-5.3%	36.1%
Gross Margin	17.4%	19.8%	23.6%	26.8%	28.5%
Adj. EBIT Margin	1.3%	3.6%	10.1%	14.6%	17.1%
Net Margin	4.4%	12.8%	14.9%	12.1%	14.1%
Net Debt/Equity	-70.6%	-62.8%	-78.3%	-121.3%	-168.6%
Return on Capital	18.3%	47.1%	42.4%	35.7%	35.8%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF Yield (%)	0.6%	1.2%	2.2%	2.4%	3.1%
Dividend yield (%)	0.0%	0.0%	0.2%	0.2%	0.2%
EV/EBITDA(x)	40.8	64.3	47.2	31.3	23.0
P/E (x)	58.9	36.9	53.0	46.4	33.4
P/B (x)	9.5	15.8	19.2	14.9	11.5
P/S (x)	2.6	4.7	6.6	5.7	5.0

Volatility

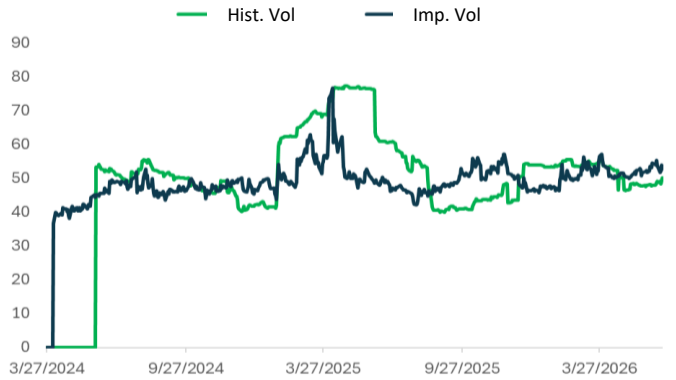

Valuation

Revenue breakdown

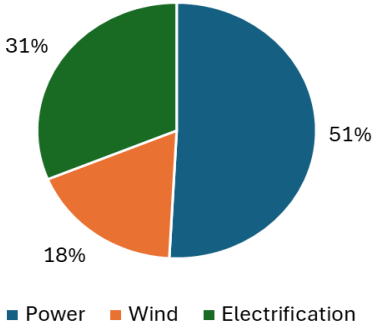

US Equities: Eaton (ETN)

1. Investment Thesis

Eaton is benefiting from Data Center and grid electrification tailwinds; its Electrical Americas (EA) segment's LTM organic orders surged. Organic orders +60% YoY; backlog +44% YoY and +9.8% QoQ; quarter-end backlog reached USD 4.9 billion. The company is carving out low-growth units for USD 5.1 billion (8.3x 2026E EBITDA) to focus on high-GM Electrical and Aerospace. In April, Eaton implemented across-the-board price increases to hedge inflation; segment margin is expected to climb from 26% to 30%+ by year-end.

2. Valuation Summary

Consensus target price range is USD 459. Some bullish views base on 2027E EPS USD 15.5, assigning c.29x forward P/E; the premium mainly reflects high visibility of core electrification order growth. While near-term margin "growing pains" warrant caution, earnings quality should improve as low-margin units are separated.

3. Share Price Catalysts

Near-term catalysts include: First, margin recovery in Electrical Americas (EA) in 2H, especially the impact of price increases; Second, spinning off the Mobility division to merge with another manufacturer, expected to complete in Q1 2027, returning USD 1.1 billion cash and delivering USD 250 million synergies; Third, new 800VDC products starting to book orders by end-2026.

4. Investment Risks

Key risks include: First, capacity expansion to meet strong orders may face execution bottlenecks during the ramp; Second, commodity inflation overshooting expectations could weaken the offset from pricing power; Finally, before the Vehicle business spin-off completes, cyclical slowdowns in the auto market and non-residential construction may pose near-term pressure.

5. Company Profile

Eaton is a global leader in power management and control systems, providing efficient, safe, and sustainable power management solutions. The company operates across Electrical Americas, Electrical Global, Aerospace, Vehicle, and eMobility. Its products and technologies serve end markets such as Data Centers and the power grid, accelerating a shift toward high-growth electrification.

Financials (USD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	24.9	27.4	31.9	35.2	38.4
Adj. EBIT (100 Mn)	4.6	5.2	6.4	7.5	8.4
Adj. EPS (USD)	10.8	12.1	13.3	15.5	17.9
OCF (100 Mn)	4.3	4.5	5.0	6.3	7.3
FCF (100 Mn)	3.5	3.6	3.9	5.2	6.1
Margins and growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue growth	7.3%	10.3%	16.3%	10.4%	9.1%
Adj. EPS growth	18.4%	11.8%	10.1%	16.9%	15.0%
Gross margin	38.2%	37.6%	37.6%	38.8%	39.0%
Adj. EBIT margin	18.6%	19.0%	20.1%	21.4%	21.8%
Net margin	15.3%	14.9%	16.2%	17.5%	18.2%
Net debt/Equity	42.7%	50.8%	86.4%	79.5%	72.2%
Return on capital	20.2%	21.6%	23.0%	24.2%	24.8%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	2.1%	2.2%	2.4%	3.2%	3.7%
Dividend yield (%)	0.3%	0.3%	1.0%	1.1%	1.2%
EV/EBITDA(x)	23.0	19.2	24.3	20.9	18.8
P/E (x)	34.9	30.5	31.7	27.2	23.6
P/B (x)	7.1	6.4	7.3	6.4	5.7
P/S (x)	5.3	4.5	5.1	4.6	4.3

Volatility

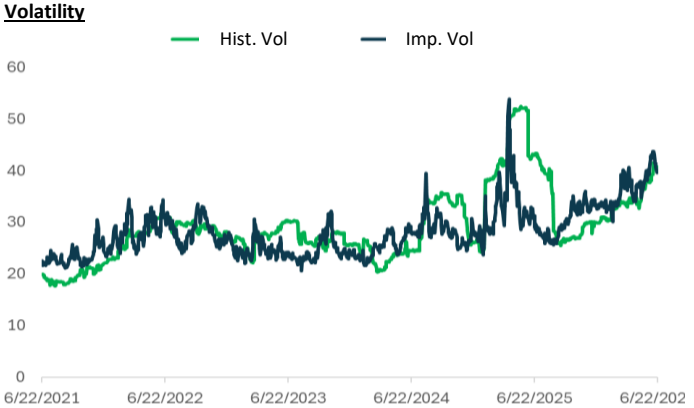

Valuation

Revenue breakdown

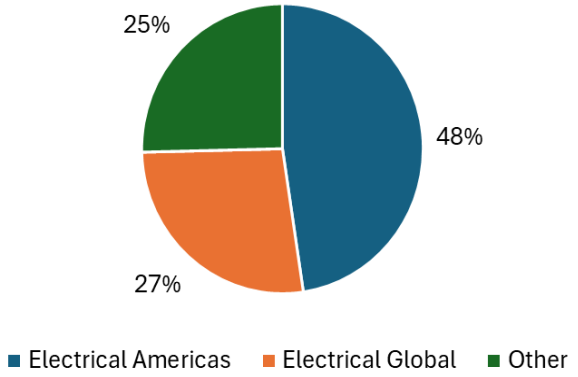

US Equities: Vertiv (VRT)

1. Investment Thesis

As AI demand surges, global data center installed capacity is now forecast at 140 GW (CAGR 15%), while TAM per MW rises to USD 3.25M-3.75M. Accordingly, the company raised its 2025-30 organic revenue CAGR guidance to 20-22% and lifted its 2030 operating margin target to above 27% (incremental margin 32%), driving 2030 base EPS to USD 14.14. Coupled with USD 28 billion of deployable capital ahead (USD 24 billion earmarked for M&A), these strong metrics underscore its robust long-term growth potential.

2. Valuation Summary

The current share price is around USD 322-328. Street target prices cluster at USD 381. The premium versus peers reasonably reflects generative AI's accelerating catalyst for the data center market and the company's technology leadership driving high earnings visibility.

3. Share Price Catalysts

Key catalysts include faster rollout of integrated prefabricated platforms such as "OneCore" that can cut deployment time by 50%; high-density compute pushing liquid cooling mix toward 90-95% to leverage full-stack cooling advantages; and greater aftermarket service needs for high-complexity equipment, accelerating service revenue growth to above 20% and sustaining improvements in operations and earnings.

4. Investment Risks

Risks to monitor: (1) constrained power availability, permitting delays, and labor shortages may push out data center completions; (2) tariffs, inflation, and supply chain challenges may erode margins; (3) softness in telecom (over 20% of revenue) and macro volatility in China (about 10-13% of revenue); and (4) a high valuation multiple lowers tolerance for any execution missteps.

5. Company Profile

The company is a global leader in critical digital infrastructure and continuity solutions, serving data centers, telecom networks, and commercial/industrial environments. It holds technology leadership in power management (e.g., UPS) and thermal management (liquid and air cooling), offering end-to-end solutions from stand-alone components and prefabricated modular platforms to monitoring software and after-sales maintenance, positioning it as a core enabler of global AI compute infrastructure buildout.

Financials (USD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	8.0	10.2	13.9	17.8	21.5
Adj. EBIT (100 Mn)	1.4	1.8	3.2	4.4	5.6
Adj. EPS (USD)	2.9	4.2	6.5	8.7	11.1
OCF (100 Mn)	1.3	2.1	2.5	3.6	4.4
FCF (100 Mn)	1.2	1.9	2.3	3.0	3.7
Margins and Growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue Growth	16.7%	27.7%	35.6%	28.2%	20.7%
Adj. EPS Growth	139.5%	47.4%	54.0%	34.0%	28.1%
Gross Margin	36.6%	36.3%	38.5%	39.3%	39.7%
Adj. EBIT margin	17.1%	17.9%	23.3%	24.6%	25.9%
Net margin	6.2%	13.0%	18.3%	17.4%	17.9%
Net debt/equity	78.6%	35.5%	-6.2%	-57.0%	-121.7%
Return on capital	22.3%	41.8%	49.3%	50.9%	47.4%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	0.9%	1.5%	1.8%	2.3%	2.9%
Dividend yield (%)	0.0%	0.0%	0.1%	0.1%	0.1%
EV/EBITDA(x)	26.3	27.7	38.0	28.2	22.1
P/E (x)	88.8	47.5	51.5	38.4	30.0
P/B (x)	17.8	15.7	21.7	15.2	10.8
P/S (x)	5.3	6.0	9.2	7.2	6.0

Volatility

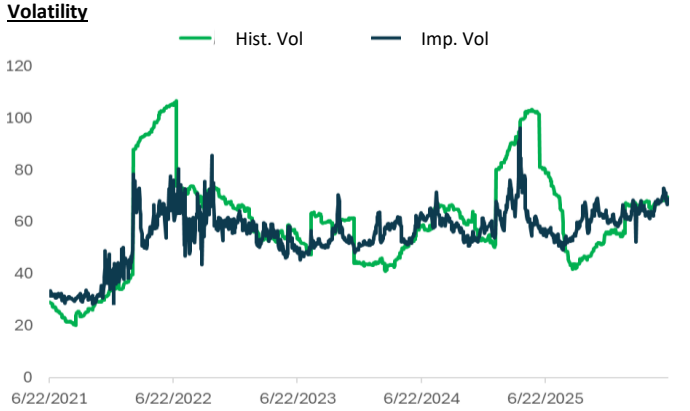

Valuation

Revenue breakdown & OPM

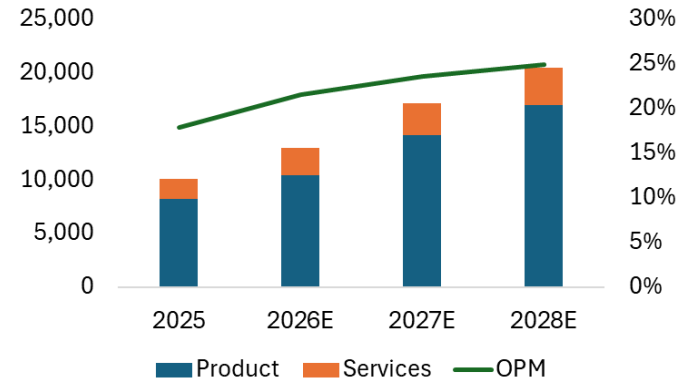

Taiwan equities: AVC (3017 TT)

1. Investment Thesis

As the AI liquid-cooling leader, AVC has strong long-term growth drivers. Product mix continues to improve, with server revenue mix rising from 54% in 4Q25 to 66% in 1Q26, lifting gross margin and operating margin to 29.8% and 24.5%. We expect liquid-cooling to reach 35-40% of revenue in 1H26, well above 15% in 2025, and surpass 50% in 2027. From 2H26, multiple marquee customer programs ramp to volume and become the core growth engine, including AMD Helios rack cold plates (2027 shipments of 5,000-8,000 racks), AWS Trainium3, Google TPU and NVIDIA VR200, driving EPS to surge to TWD 94.49 in 2026 and TWD 132.70 in 2027.

2. Valuation Summary

Applying 24x 2027E EPS yields a 12-month target price of TWD 3,360. With liquid-cooling shipments accelerating in 2H and mix improving, a valuation re-rating is likely; maintain Add.

3. Share Price Catalysts

Near-term catalysts include significant capacity expansion (cold-plate monthly capacity to rise to 1 million units) and small-batch VR200 cold-plate shipments by end-2Q26. Medium to long term, AWS Trainium3 liquid-cooling and Google TPU programs will ramp to volume from 4Q26 and contribute revenue; along with 2027 shipments of AMD Helios rack cold plates and internal manifolds and other high-value new products, this will drive revenue and profits to strengthen sequentially from 2H26, significantly bolstering confidence in its liquid-cooling leadership.

4. Investment Risks (Investment Risks)

Key risks include: (i) delays in AI server upgrades or slower pull-ins from CSP customers (e.g., deferred GB300 demand) affecting quarterly growth; (ii) NTD appreciation diluting GM and causing FX losses; (iii) rising raw material prices (copper, aluminum) increasing costs; (iv) technology shifts or intensified competition compressing ASPs for thermal modules and cold plates.

5. Company Profile (Company Profile)

Founded in 1991, the company is a global leader in thermal solutions, offering thermal modules, fans, and liquid cooling systems. In 1Q26, thermal accounted for 64% of revenue, racks 20%, and foldable display hinges (Fushida) 9%. The company has transitioned into an integrated supplier of thermal and mechanical components, leveraging high in-house rates for cold plates, manifolds, and strong system design capabilities to sustain industry leadership.

Financials (TWD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	71.8	139.6	220.4	298.5	341.5
Adj. EBIT (100 Mn)	10.8	27.6	55.1	78.6	92.6
Adj. EPS (TWD)	21.2	49.2	94.5	136.4	160.8
OCF (100 Mn)	9.6	39.4	27.5	67.9	60.0
FCF (100 Mn)	5.0	31.9	-4.9	24.2	44.7
Margins and growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue growth	21.2%	94.6%	57.8%	35.5%	14.4%
Adj. EPS growth	54.7%	131.8%	92.3%	44.3%	17.9%
Gross margin	23.5%	25.8%	30.4%	31.3%	30.7%
Adj. EBIT margin	15.1%	19.7%	25.0%	26.3%	27.1%
Net margin	11.4%	13.7%	16.9%	17.8%	17.6%
Net debt/Equity	-24.0%	-50.1%	-96.4%	-105.8%	-174.7%
Return on capital	32.1%	52.2%	53.9%	47.0%	40.0%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	0.5%	3.4%	-0.5%	2.5%	4.7%
Dividend yield (%)	-	-	1.3%	2.0%	2.2%
EV/EBITDA(x)	17.9	18.5	15.5	11.1	9.2
P/E (x)	30.2	31.3	25.6	17.7	15.1
P/B (x)	8.4	13.2	12.9	9.0	5.7
P/S (x)	3.3	4.2	4.3	3.2	2.8

Volatility

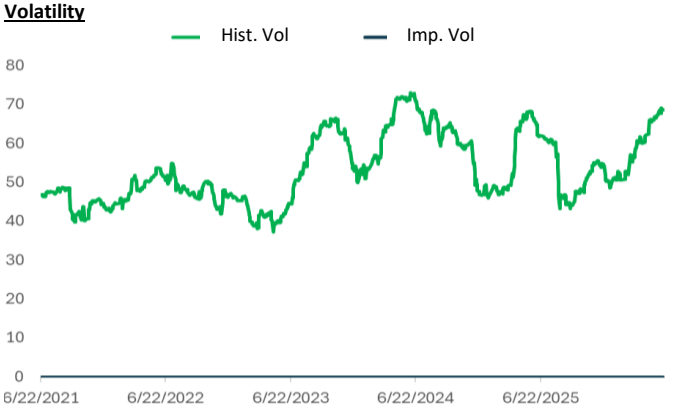

Valuation

Valuation Breakdown

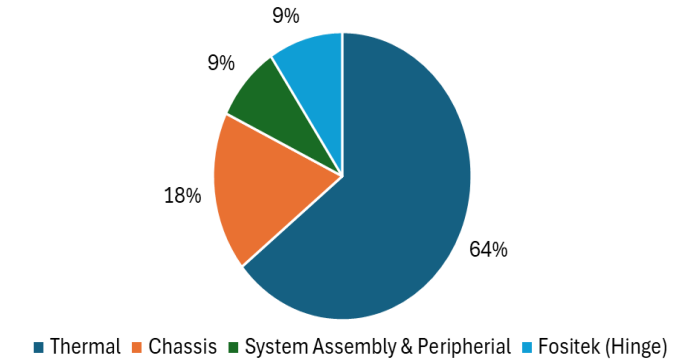

US: Intel (INTC)

1. Investment Thesis

Benefiting from robust demand for Agentic AI, CPUs are regaining a central role in task orchestration and tool invocation, narrowing the data center CPU-to-GPU ratio from the traditional 1:4–1:8 to 1:1–1:2. We expect server CPU shipments to post double-digit growth in 2026. Meanwhile, custom ASIC revenue has already surpassed USD 1 billion annually. On manufacturing, core 18A process yields reached target six months ahead of plan, and the Terafab fab project with Musk is slated to inject USD 20–25 billion in Capex. Process leadership, ASIC expansion, and strong progress on Terafab underpin a high-growth runway for its foundry and product businesses.

2. Valuation Summary

The market predominantly uses SOTP, implying a fair target price of USD 99.5. The model values the product businesses (CCG and DCAI) at 23x EV/EBIT in line with peers; values foundry capacity at USD 193–255 per 1,000 equivalent wafers; and deducts non-controlling interests from JVs such as Brookfield.

3. Share Price Catalysts

Key near-term catalysts include releasing the 1.0 PDK for 14A by end-2026, expected to catalyze tech giants to sign formal foundry commitments in 2H26–1H27; volume shipments of Panther Lake processors on 18A; finalization of the partnership structure for Musk's Terafab project; and the foundry business reaching operating quarterly breakeven by end-2027—all driving multiple re-rating and bolstering investor confidence.

4. Investment Risks

Key risks include: (1) sluggish PC end-demand recovery, with global PC TAM in 2026 expected to decline 5% to double digits, pressuring near-term profitability; (2) heavy Capex burden from the foundry transition, with 2026 Capex remaining elevated at USD 18 billion, weighing on FCF; (3) ongoing share gains by competitors in server CPUs, and profit sharing to non-controlling interests from early SCIP JV arrangements diluting EPS.

5. Company Profile

The company is a global semiconductor leader with about 60% share in x86 microprocessors, dominating the PC and data center CPU markets. It is accelerating its transformation into a vertically integrated (IDM 2.0) foundry services provider and offers leading 3D advanced packaging. It also owns Mobileye and Altera, building a broad ecosystem spanning compute, manufacturing, and smart mobility.

Financials (USD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	53.1	52.9	58.7	65.4	72.9
Adj. EBIT (100 Mn)	-2.5	2.9	7.1	10.0	13.7
Adj. EPS (USD)	-0.1	0.4	1.1	1.6	2.4
OCF (100 Mn)	8.3	9.7	16.1	21.5	26.2
FCF (100 Mn)	-15.7	-4.9	0.8	5.0	8.0
Margins & growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue growth	-2.1%	-0.5%	11.0%	11.3%	11.5%
Adj. EPS growth	-	-	162.4%	42.7%	51.3%
Gross margin	36.0%	36.7%	40.4%	42.9%	46.4%
Adj. EBIT margin	-4.8%	5.5%	12.2%	15.3%	18.8%
Net margin	-1.1%	3.7%	9.6%	12.8%	17.8%
Net debt/equity	21.9%	0.8%	3.8%	-4.6%	-19.0%
Return on capital	-18.3%	-0.3%	3.6%	6.5%	8.2%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	-2.3%	-0.7%	0.1%	0.7%	1.2%
Dividend yield (%)	0.0%	0.0%	0.0%	0.0%	0.0%
EV/EBITDA(x)	19.1	15.8	38.8	30.8	26.3
P/E (x)	-	-	121.6	85.2	56.3
P/B (x)	0.9	1.6	5.6	5.3	4.4
P/S (x)	1.6	3.1	11.5	10.3	9.2

Volatility

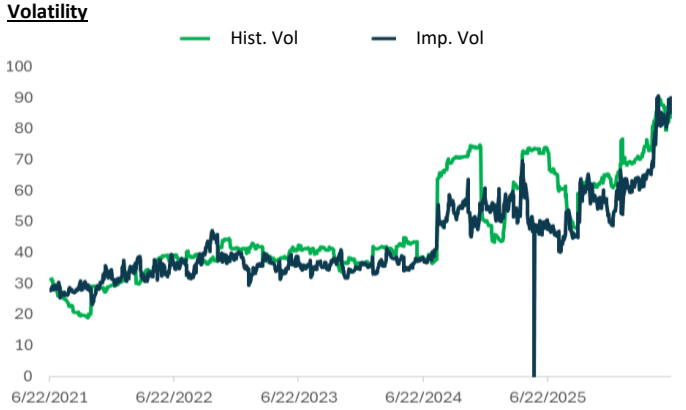

Valuation

Revenue breakdown

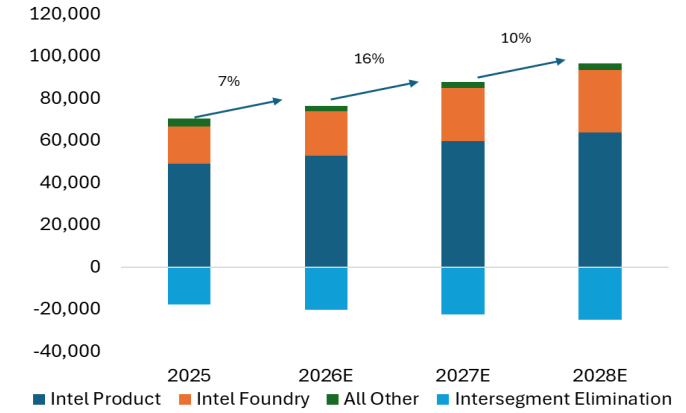

US: Dell Technologies (DELL)

Investment Thesis

First-quarter results surged on AI and traditional servers. Quarterly revenue reached USD 43.8 billion (+88% YoY), well above guidance; EPS was USD 4.86. AI server revenue in the first quarter jumped to USD 16.1 billion; quarterly orders were USD 24.4 billion; ending backlog hit a record USD 51.3 billion; customer count topped 5,000. Meanwhile, traditional server revenue rose 92% YoY to USD 8.5 billion on the 14G-to-18G enterprise upgrade cycle. Backed by roughly 50 customized liquid-cooling configurations, the company raised its full-year AI revenue outlook to USD 60.0 billion, with a clearly visible growth moat.

Valuation Summary

Street target price raised to USD 485, implying ~22x forward P/E. Versus peers at 11–32x (median 19x), the company's strong AI leadership and 100% FCF conversion warrant a reasonable premium; profitability improvement should drive multiple expansion.

Share Price Catalysts

Key catalysts include: rolling out next-gen liquid-cooled servers next year, expanding ASPs and margins; higher penetration of in-house, high-margin storage products, improving the profit mix; and the rise of edge-side Agentic AI applications, triggering upgrade cycles for high-end PCs and traditional servers. In addition, strong cost pass-through capability should safeguard GM amid component price inflation.

Investment Risks

Key risks include: severe component tightness, with DRAM and CPU lead times up to one year—price increases may constrain shipments and squeeze margins; a tepid PC market recovery, with potential H2 price wars driving segment margins back toward ~6%; and competition from white-box ODMs and cloud service providers, which could erode share and pricing power in traditional servers and storage.

Company Profile

Founded in 1984, the company is a global leader in IT hardware and personal computing solutions. It operates two segments: Infrastructure (ISG, covering AI and traditional servers, storage, and networking) and Client (CSG, covering PCs). With a world-class supply chain and customized engineering/deployment capabilities, the company holds a leading position in the current digital transformation and AI build-out.

Financials (USD)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	95.6	113.5	170.6	194.1	211.4
Adj. EBIT (100 Mn)	8.5	10.0	15.9	18.2	20.7
Adj. EPS (USD)	8.1	10.3	17.9	21.7	25.6
OCF (100 Mn)	4.5	11.2	17.2	18.5	21.1
FCF (100 Mn)	1.9	8.6	12.7	14.6	16.9
Margins and Growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue Growth	8.1%	18.8%	50.2%	13.8%	8.9%
Adj. EPS Growth	14.2%	26.5%	74.0%	20.8%	18.2%
Gross Margin	22.4%	20.2%	17.8%	17.5%	17.6%
Adj. EBIT Margin	8.9%	8.8%	9.3%	9.4%	9.8%
Net profit margin	4.8%	5.2%	7.0%	7.1%	7.6%
Net debt/equity	-	-	-677.5%	-457.1%	-271.1%
Return on capital	-	-	-	-	-
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	0.7%	3.2%	4.8%	5.5%	6.4%
Dividend yield (%)	0.2%	0.2%	0.6%	0.7%	0.7%
EV/EBITDA(x)	9.0	7.7	15.2	13.3	11.5
P/E (x)	16.2	13.2	22.8	18.9	16.0
P/B (x)	-	-	247.3	47.6	46.6
P/S (x)	0.8	0.7	1.6	1.4	1.3

Volatility

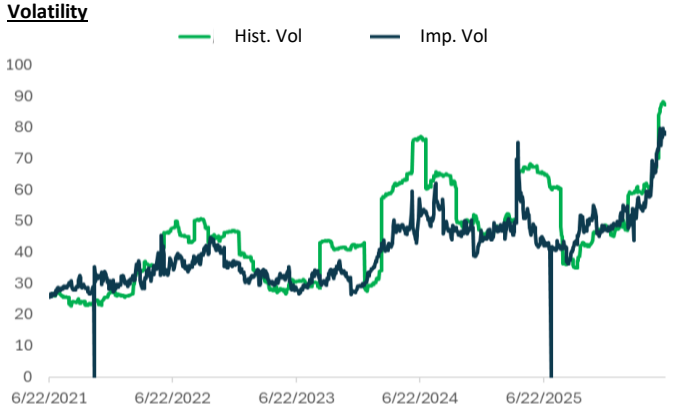

Valuation

Orderflow

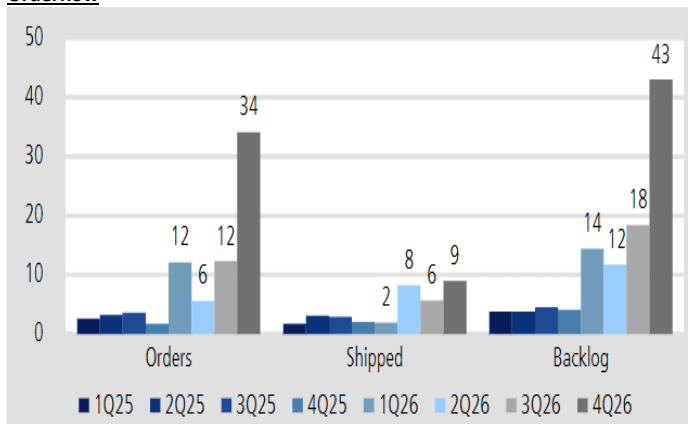

Hong Kong/China ADRs: Tencent (00700 HK)

1. Investment Thesis (Investment Thesis)
Core businesses demonstrate exceptional resilience. In 1Q26, non-IFRS EPS rose 12% YoY, well ahead of market expectations of a 5–8% increase. Non-IFRS net profit reached ~RMB 67.9 billion in the quarter, up 10.7% YoY. On AI investment, the split P&L established a clear annual framework; quarterly AI opex is ~RMB 8.8 billion (annualized ~RMB 36.0 billion), alleviating concerns about unchecked R&D expansion. Core engines are operating solidly, with marketing services growth accelerating to 20% and business services sustaining 20% growth. The in-house Hunyuan model is deeply integrated into 131 internal products, with 1Q token consumption more than 10x that of the prior Hunyuan 2. With robust FCF of RMB 57.0 billion and a net cash position, the company can strongly fund model iteration and self-financed R&D, making it a scarce, high-quality compounding earnings asset in the market.

2. Valuation Summary (Valuation Summary)
Shares trade at only 12–13x forward P/E, near a five-year trough, offering strong defensiveness and downside support. The market's prevailing approach is SOTP, implying a fair target price range of HKD 687. As AI investment visibility improves and EPS expectations are revised up, the undervalued AI option value should be unlocked, driving multiple re-rating.

3. Share Price Catalysts (Share Price Catalysts)

Key catalysts include testing and official launch of the WeChat AI agent (Weixin Agent), notably its deep integration with the WeChat Mini Program ecosystem; the first batch of partners such as Meituan and Trip.com have joined trials, which should materially enhance user experience. In addition, the release of next-gen models (e.g., HY4), acceleration in gross billings for key games domestically and overseas, and the company's robust buybacks of about HKD 500 million per day should continue to bolster market confidence.

4. Investment Risks (Investment Risks)
Key risks to watch: (1) Capex expansion—quarterly Capex is rising sharply; an overly fast ramp could reignite concerns over discipline in AI spend; (2) macro slowdown weighing on advertising and fintech growth; (3) during the early rollout of new AI products, elevated inference compute costs may dilute near-term GM.

5. Company Profile
Tencent (00700 HK) is China's leading integrated internet services provider, built on a vast social-network ecosystem anchored by WeChat and QQ; WeChat Mini Programs connect millions of third-party merchants, creating an unrivaled private-domain monetization channel. The company is a market leader in online games, online advertising, and fintech & cloud services, and is accelerating R&D and integration of multiple AI productivity tools for developers and office workflows.

Financials (RMB)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	660.3	751.8	829.1	907.3	994.5
Adj. EBIT (100 Mn)	208.1	241.6	271.4	301.2	336.4
Adj. EPS (RMB)	23.5	27.9	30.2	33.5	37.3
OCF (100 Mn)	264.5	310.9	321.4	363.1	417.9
FCF (100 Mn)	194.9	223.4	189.1	231	276.7
Margins and growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue growth	8.4%	13.9%	10.3%	9.4%	9.6%
Adj. EPS growth	44.0%	18.6%	8.4%	11.0%	11.1%
Gross Margin	52.9%	56.2%	56.9%	57.5%	58.0%
Adj. EBIT margin	31.5%	32.1%	32.7%	33.2%	33.8%
Net margin	29.4%	29.9%	33.5%	33.6%	34.2%
Net debt/equity	1.4%	-1.6%	10.6%	4.8%	-2.7%
Return on capital	21.8%	21.1%	19.4%	18.8%	18.2%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	5.0%	5.7%	4.8%	5.9%	7.0%
Dividend yield (%)	1.1%	1.1%	1.2%	1.3%	1.5%
EV/EBITDA(x)	14.0	16.1	10.2	9.2	8.2
P/E (x)	19.1	22.3	12.4	11.2	10.1
P/B (x)	3.7	4.2	2.5	2.2	2.0
P/S (x)	5.5	6.5	4.1	3.8	3.4

Volatility

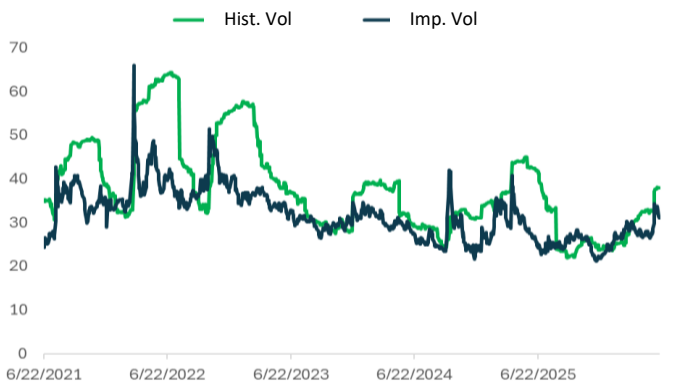

Valuation

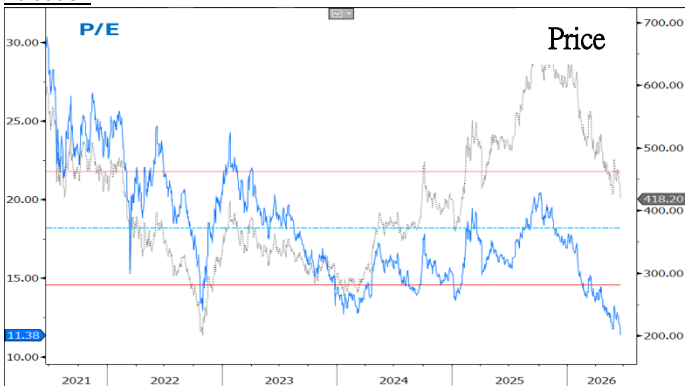

Short-selling activity in Tencent (00700 HK) is gradually easing

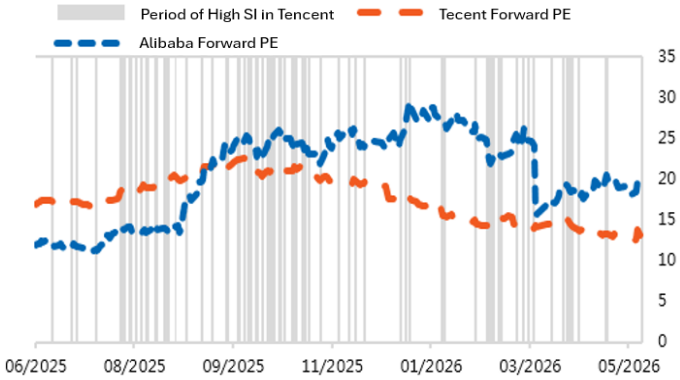

Hong Kong/China ADRs: Alibaba (BABA / 09988 HK)

1. Investment Thesis

Core businesses demonstrate strong resilience; Alibaba Cloud has emerged as China's leading AI+cloud champion, with embedded value not fully reflected in the current share price. Cloud external revenue YoY growth improved to 40%, with AI-related annualized revenue reaching RMB 36 billion, delivering triple-digit growth for 11 consecutive quarters. Standalone MaaS ARR has surpassed RMB 8 billion and is expected to exceed RMB 30 billion by year-end. Cloud EBITA margin rose 110 bps YoY to 9.1% and is poised to reach double digits in the next 1–2 quarters. Its T-Head chip ecosystem is robust; flagship Zhenwu AI chip shipments have reached 560,000 units, serving 400+ customers; the new M890 chip markedly improves AI inference efficiency. Core e-commerce revenue grew 6%, remaining solid. These data points underscore substantial AI and cloud re-rating potential with powerful structural growth drivers.

2. Valuation Summary

The current share price mainly reflects the legacy e-commerce business, with the sizable option value of cloud and AI underappreciated by the market. The market largely applies SOTP, implying a reasonable target-price range of HKD 180. This assigns 10 P/E to core e-commerce profit and 6–7x P/S to cloud. As cloud EBITA margin returns to double digits in the next 1–2 quarters, implied AI re-rating should accelerate.

3. Share Price Catalysts

Key catalysts include: cloud EBITA margin returning to double digits within the next 1–2 quarters; launch of the next-gen flagship model Qwen 3.7-Max and in-house T-Head Zhenwu M890 chip; accelerated narrowing of losses in new businesses such as on-demand retail, with a commitment to halve losses by FY2027/28 and reach profitability in FY2029; and continued capital return via sizable share buybacks and a stable dividend, driving a valuation inflection.

4. Investment Risks

Key risks include: elevated AI Capex potentially pressuring margins near term alongside ongoing compute-chip supply constraints; intense competition in domestic e-commerce and local services may delay monetization improvement; macro slowdown could soften advertising and cloud demand; plus regulatory and international trade uncertainties.

5. Company Profile

Through Taotian Group and International Retail, the company has built a powerful commerce ecosystem, providing technology infrastructure and marketing platforms to millions of merchants. It is also China's largest AI and cloud operator, rapidly pivoting to the Agent Era, empowering enterprise and public-sector intelligent transformation via the Bailian platform and Qwen foundation models.

Financials (RMB)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	996.3	1023.7	1123.2	1256.9	1415.5
Adj. EBIT (100 Mn)	140.9	50.2	92.0	139.4	193.3
Adj. EPS (RMB)	65.4	26.8	43.4	61.6	81.5
OCF (100 Mn)	163.5	76.2	196.7	211.1	307.7
FCF (100 Mn)	77.5	-49.9	20.9	46.8	98.6
Margins and growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue growth	5.9%	2.7%	9.7%	11.9%	12.6%
Adj. EPS growth	5.1%	-59.0%	62.1%	41.8%	32.4%
Gross margin	40.0%	39.8%	39.4%	40.2%	40.9%
Adj. EBIT margin	14.1%	4.9%	8.2%	11.1%	13.7%
Net margin	13.1%	10.1%	9.8%	12.0%	13.7%
Net debt/Equity	-49.2%	-44.0%	-2.4%	-3.4%	-13.0%
Return on capital	13.0%	10.2%	9.0%	11.3%	14.1%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	30.2%	-19.4%	8.1%	18.2%	38.4%
Dividend yield (%)	-	-	6.7%	8.9%	7.8%
EV/EBITDA(x)	8.8	34.2	8.3	6.1	4.6
P/E (x)	17.9	19.7	16.7	11.8	8.9
P/B (x)	2.3	2.0	1.7	1.5	1.3
P/S (x)	2.3	2.0	1.6	1.4	1.2

Volatility

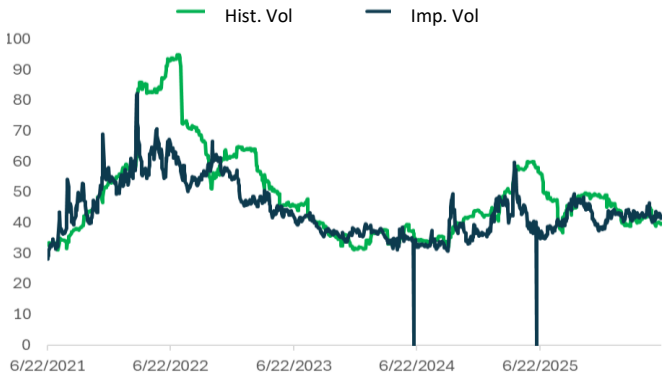

Valuation

AliCloud Performance

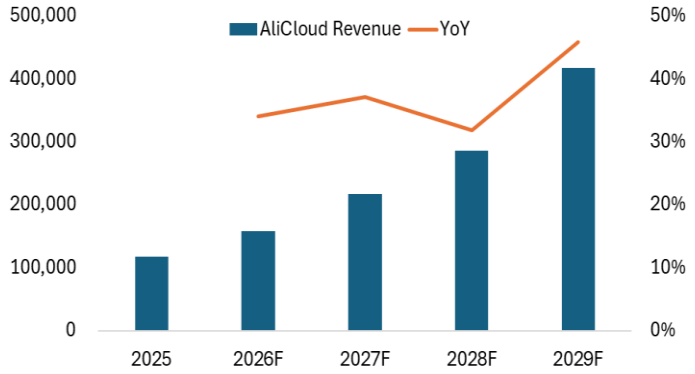

Hong Kong/China ADRs: Baidu (BIDU / 09888 HK)

1. Investment rationale (Investment Thesis)

Successfully transformed into an AI infrastructure flagship. In 1Q26, core AI revenue jumped 49% to RMB 13.6 billion, accounting for 52% of core revenue for the first time. Intelligent Cloud infrastructure surged 79% to RMB 8.8 billion, with GPU cloud growth exploding 184%, driving core adjusted operating profit to RMB 4.0 billion. In-house Kunlunxin is set to ship 116k units in 2025, with 2026 revenue expected to reach RMB 8.7 billion. In addition, Apollo Go quarterly orders rose 120% to 3.2 million; the "Cloud + Chips + Autonomy" full-stack in-house ecosystem rests on a very solid data foundation.

2. Valuation Summary

The current share price is at historical lows. The market applies SOTP, with a fair target price of HKD 171. The per-share value of just AI Cloud (3–6x P/S), Kunlunxin chips (15–35x P/S) and Autonomous Driving AI businesses already exceeds the current share price—implying the market is assigning zero consideration to legacy search, the iQIYI (IQ) stake and substantial net cash, resulting in a very high margin of safety.

3. Share Price Catalysts

Key rerating catalysts include: (1) in-house AI chip Kunlunxin pursuing "A+H" dual listings (STAR Market and HKEX), potentially unlocking up to HKD 100 billion in valuation; (2) Apollo Go, the world's largest autonomous ride-hailing service, doubling orders YoY, with overseas and Hong Kong cross-district trials accelerating, pointing to a clear path to profitability; (3) ERNIE 5.1 iteration and DuMate rollout driving MaaS token consumption and cloud margin expansion.

4. Investment Risks

Key risks include: first, ongoing pressure on legacy search advertising revenue (1Q26 legacy search down 29% YoY), share losses to competitors and a slower-than-expected recovery; second, to match high AI growth, 1Q26 Capex tripled—elevated compute equipment purchases and R&D reinvestment could dilute near-term GM; finally, a Kunlunxin spin-off listing may face a holdco discount, and global semiconductor regulation remains uncertain.

5. Company Profile

The company is China's leading AI and search-engine champion. Since 2012 it has focused on deep learning, building a vertically integrated, full-stack AI ecosystem spanning in-house AI chips (Kunlunxin), the underlying framework (PaddlePaddle), LLMs (ERNIE) and end-user agents (DuMate, Miaoda, etc.). Its Intelligent Cloud has become the leading domestic GPU cloud compute platform, while Apollo Go is also the world's largest autonomous driving service provider, spearheading the IoT and agent era.

Financials (RMB)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	133.1	129.1	132.5	142.7	156.5
Adj. EBIT (100 Mn)	21.3	-5.8	14.0	17.0	20.8
Adj. EPS (RMB)	76.9	53.4	50.4	59.4	73.6
OCF (100 Mn)	21.2	-3.0	27.3	30.7	44.6
FCF (100 Mn)	13.1	-15.1	7.6	14.8	16.2
Margins and growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue growth	-1.1%	-3.0%	2.7%	7.7%	9.7%
Adj. EPS growth	39.4%	-30.5%	-5.7%	18.0%	23.8%
Gross margin	50.3%	43.9%	40.9%	42.4%	43.1%
Adj. EBIT margin	16.0%	-4.5%	10.6%	11.9%	13.3%
Net margin	17.8%	4.3%	13.5%	14.3%	15.7%
Net debt/Equity	-51.7%	-49.0%	0.9%	0.1%	-7.7%
Return on capital	9.1%	1.8%	5.1%	5.7%	6.7%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	34.4%	-39.7%	19.9%	38.8%	42.5%
Dividend yield (%)	0.0%	0.0%	0.0%	0.0%	0.0%
EV/EBITDA(x)	2.9	8.6	6.1	5.1	4.4
P/E (x)	9.3	67.2	15.0	12.7	10.3
P/B (x)	0.8	1.2	0.9	0.8	0.8
P/S (x)	1.6	2.4	1.9	1.8	1.6

Volatility

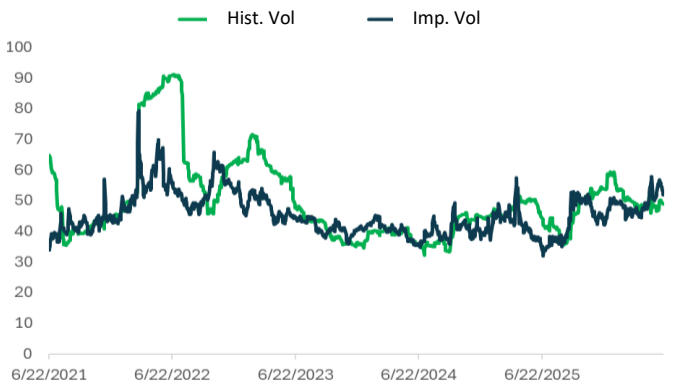

Valuation

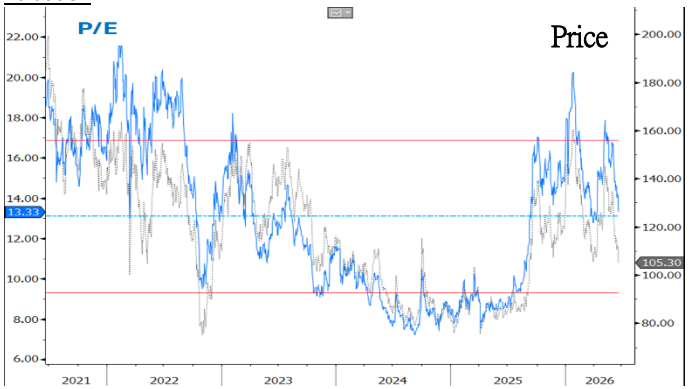

CAPEX

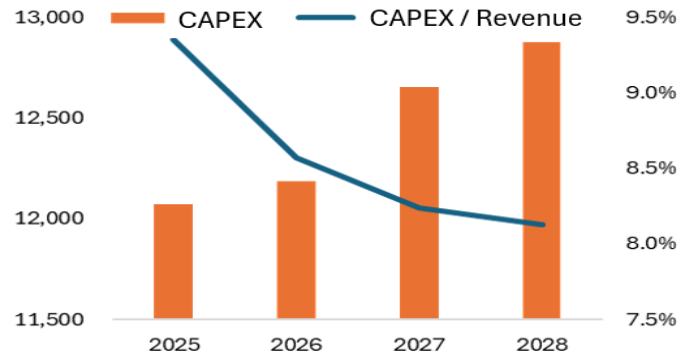

A/H shares: VGT (02476 HK)

1. Investment Rationale (Investment Thesis)
The global AI boom is driving demand for high-end PCB, with ultra-high-layer PCBs and high-end HDI expected to grow at 15% and 14% CAGR globally through 2029. As a key beneficiary, the company ranks No.1 worldwide in AI/HPC PCB (market share 13.8%), ultra-high-layer PCB (16.2%), and high-end HDI (44.6%). Benefiting from deep strategic partnerships with global tech leaders on the NVIDIA (NVDA) Rubin platform and Alphabet (GOOGL) ASIC projects, the company's product mix has improved significantly, driving 2025 net profit up 274% to RMB 4.3 billion. We expect a net profit CAGR of 81% over 2025–2028, and with a capacity expansion plan of up to RMB 18.0 billion in 2026, strong order visibility should propel earnings into a hypergrowth phase.

2. Valuation Summary (Valuation Summary)
The company is dual-listed in A-shares and H-shares. Our A-share target price is RMB 442, implying 25x 2027E P/E, in line with the 3-year historical average, reflecting robust growth potential driven by AI demand. Our H-share target price is HKD 539, based on 34x 2026–2027E average P/E.

3. Share Price Catalysts (Share Price Catalysts)

Key catalysts include: first, hardware platform upgrades by global tech giants and mass-production ramps in 2H will drive substantial incremental orders; second, ASIC customers such as the 9th-gen Alphabet (GOOGL) TPU will step up procurement of high-end HDI, lifting margins; third, new lines in Vietnam and Thailand are expected to come online progressively in 2026, with overseas capacity easing geopolitical risk concerns; finally, from 2H26, quarterly earnings growth should accelerate markedly, prompting a re-rating in both markets.

4. Investment Risks (Investment Risks)
Key risks to watch: first, the large 2026 capacity addition faces uncertainties around new-plant yields and ramp progress, which could pressure gross margin in the short term; second, intense industry competition and the risk of share loss if major customers qualify new suppliers; third, sharp volatility in copper and other key raw material prices—with raw materials accounting for over 60% of COGS—can squeeze margins; finally, international trade frictions, tariff policy changes, and potential geopolitical tensions.

5. Company Profile (Company Profile)
Founded in 2003, a dual-listed (Shenzhen and Hong Kong) global leader in high-end PCB manufacturing. Core businesses include ultra-high-layer-count boards, advanced HDI, and FPC, with products embedded in AI compute, automotive electronics, and smart devices. Provides global tech leaders with one-stop, highly reliable, volume manufacturing and delivery services.

Financials (RMB)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	10.7	19.3	33.8	57.6	81.7
Adj. EBIT (100 Mn)	1.3	5.1	10.5	18.6	26.5
Adj. EPS (RMB)	1.3	5.0	9.8	17.4	25.1
OCF (100 Mn)	1.3	4.5	7.4	17.4	25.0
FCF (100 Mn)	1.0	-1.7	-	-	-
Profitability and growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue growth	35.3%	79.8%	75.1%	70.6%	41.7%
Adj. EPS growth	71.8%	273.9%	95.8%	77.6%	44.0%
Gross margin	22.7%	35.2%	40.0%	41.4%	42.4%
Adj. EBIT margin	12.3%	26.6%	31.0%	32.3%	32.5%
Net margin	10.8%	22.4%	26.2%	27.7%	28.8%
Net debt/equity	62.4%	39.1%	-38.2%	-70.8%	-152.3%
Return on capital	13.9%	33.8%	28.9%	30.7%	34.6%
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	0.3%	-0.5%	-	-	-
Dividend yield (%)	-	-	0.9%	1.6%	2.5%
EV/EBITDA(x)	19.3	42.0	30.0	16.9	12.1
P/E (x)	31.4	57.5	37.2	20.9	14.6
P/B (x)	4.0	15.1	10.6	7.6	5.4
P/S (x)	3.4	12.8	10.6	6.2	4.4

Volatility

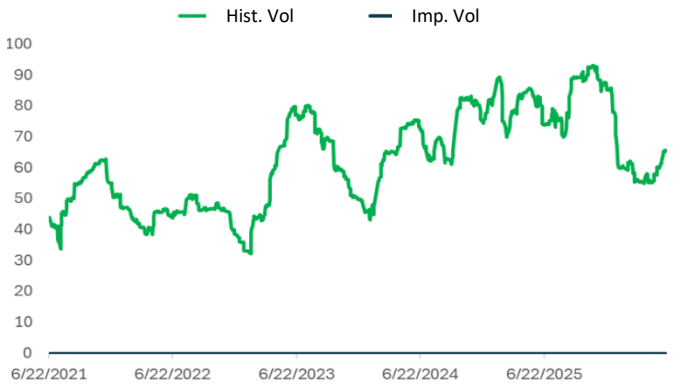

Valuation

Revenue by product type

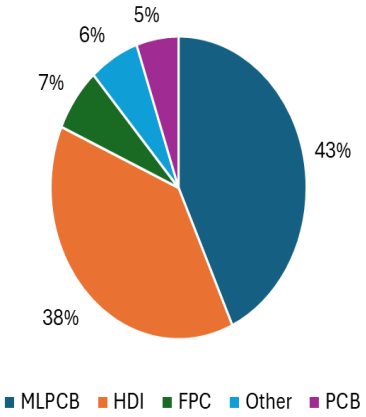

Hong Kong/China ADRs: Zhipu AI (2513 HK)

1. Investment Rationale (Investment Thesis)
Zhipu AI has top-tier frontier model R&D capabilities. The latest GLM-5.2 model's API blended pricing is 13% higher than the prior generation, and is 1.2x that of Kimi K2.7 and 4.9x that of Shenyan, demonstrating strong pricing power. ARR grew over 6x in four months to USD 250 million. Meanwhile, GLM-5.1 ranks No.1 among open-source on METR for 8-hour task completion. We expect the cloud open platform to deliver a 2025–2027 CAGR of 470%, driving a commercialization inflection.

2. Valuation Overview (Valuation Summary)
We set a TP of HKD 1,373 for Zhipu AI, based on 22x 2030E P/E, discounted at 15% WACC; the premium reflects >100% forward revenue growth. The company is expected to turn profitable in 2028 (net profit RMB 1.287 billion), rising to RMB 38.576 billion by 2030.

3. Catalysts (Share Price Catalysts)

Near-term catalysts include: next-gen trillion-parameter model GLM-5.5 slated for Aug 2026, extending its tech lead; continued progress on overseas sovereign AI platforms; and aggressive buybacks of about HKD 500 million per day, which should bolster market confidence and drive multiple re-rating.

4. Risk Factors (Investment Risks)
Risks include: geopolitics and export controls (some entities on the US Entity List could face compute access constraints); intensifying competition raising key-account churn risk, with model commoditization pressuring pricing; and sustained heavy R&D spend (2–5x revenue) plus inference compute costs, which will continue to weigh on GM and delay break-even.

5. Company Overview (Company Profile)
Zhipu AI is China's leading general-purpose foundation model and AGI pioneer, spun out of Tsinghua University and collaborating with Tsinghua's Knowledge Engineering Lab (KEG), with a top-tier team of 800+. The company offers a comprehensive MaaS platform across language, multimodal, and code models. Its business model combines on-prem/private deployments (74% mix) with high-growth cloud API services, and it has established leadership in finance, e-government, and other verticals.

Financials (RMB)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue (100 Mn)	0.3	0.7	3.2	8.5	19.8
Adj. EBIT (100 Mn)	-2.6	-3.8	-4.7	-4.9	-3.2
Adj. EPS (RMB)	-154.1	-8.1	-10.8	-10.4	-5.2
OCF (100 Mn)	-2.2	-2.2	-2.3	-1.3	2.4
FCF (100 Mn)	-2.4	-2.3	-	-	-
Margins and growth (%)	-FY2	-FY1	FY1E	FY2E	FY3E
Revenue growth	150.9%	131.9%	341.6%	166.0%	132.9%
Adj. EPS growth	195.3%	-94.7%	32.1%	-3.8%	-50.0%
Gross margin	56.3%	41.0%	35.2%	37.2%	40.5%
Adj. EBIT margin	-816.8%	-524.8%	-147.6%	-57.7%	-15.9%
Net margin	-946.3%	-648.6%	-	-	-
Net debt/equity	-	-	19.0%	33.9%	42.7%
Return on capital	-	-	-	-	-
Valuation	-FY2	-FY1	FY1E	FY2E	FY3E
FCF yield (%)	-0.2%	-0.2%	-	-	-
Dividend yield (%)	-	-	0.0%	0.0%	0.0%
EV/EBITDA(x)	-	-	-	-	-
P/E (x)	-	-	-	-	-
P/B (x)	-	-	-	-	-
P/S (x)	-	-	283.7	106.6	45.8

Volatility

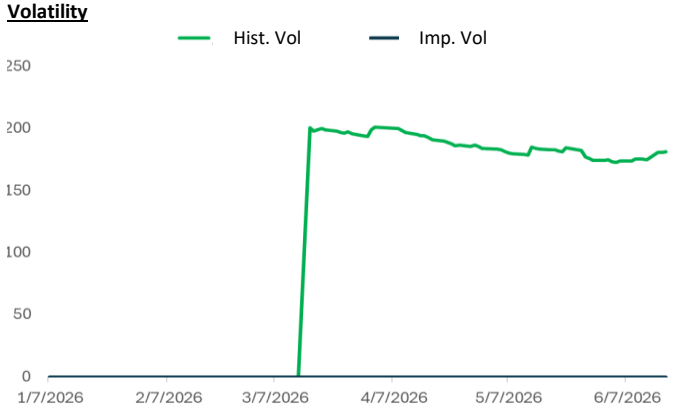

Valuation

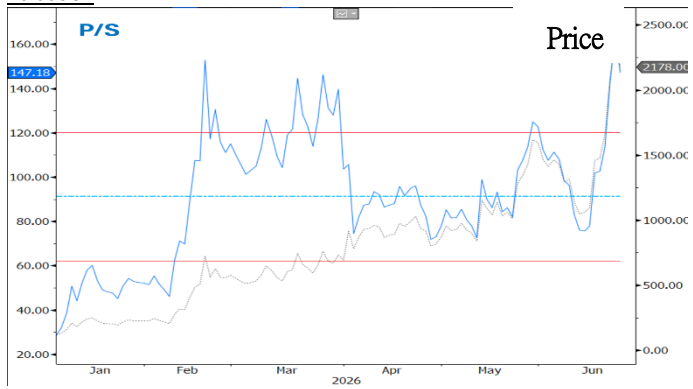

R&D Spending

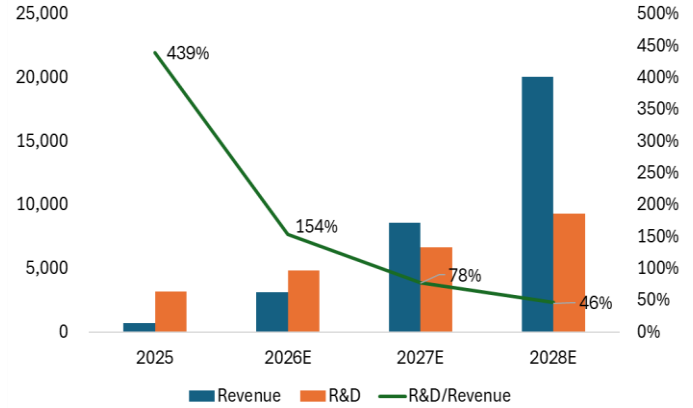

10. Glossary of Terminology

1.6T Optical Transceivers: equipment that transmits optical signals between AI compute clusters, currently migrating from 800G to 1.6T (key incremental growth drivers for Coherent (COHR) and Lumentum (LITE)).

3D VC (3D Vapor Chamber, three-dimensional vapor chamber) / Liquid Cooling: AI chips are extremely power-hungry; traditional fan-based air cooling can no longer dissipate the heat, necessitating advanced thermal solutions (a core business of Vertiv (VRT)).

AI Agent (Intelligent Agent): Not only can chat, but can also autonomously decompose tasks, call tools, and execute complex operations (e.g., automatically booking flights, writing reports, and sending emails).

AI Fatigue (AI fatigue): A sentiment that emerged in 2026, referring to users and enterprises becoming jaded by the proliferation of LLMs and wrapper apps, leading capital to concentrate in giants with absolute moats.

Applications & Eco: applications and ecosystem layer.

ASIC (Application-Specific Integrated Circuit, special-purpose integrated circuit): Chips customized for specific AI tasks, such as Alphabet (GOOGL) TPU or in-house AI chips from tech giants; more power-efficient and lower-cost than general-purpose GPUs.

ASV (AI Server Value): The value per AI server. Sell-side analysts often use it to estimate gross margin uplift for ODMs such as Foxconn Industrial Internet (601138 CH) and Quanta Computer (2382 TT).

Capex (Capital Expenditure, capital expenditure): A high-frequency term in research, specifically the money tech giants spend buying GPUs and building data centers. When CapEx rises, it implies stronger earnings for hardware stocks.

Churn Rate (attrition rate): A core metric investors use to assess C-end AI apps (e.g., MiniMax, Character.ai), defined as the share of paying users who cancel in the second month. If churn is too high, the product is mere “Novelty” rather than a must-have.

Compute (computing power): The hard currency of the AI era, often referring to the total available GPU resources.

Context Window (context window): The maximum amount of text an AI can read and retain in a single pass. For example, “200k Context” means the AI can read and understand a 200,000-word novel at once.

CoWoS (Chip-on-Wafer-on-Substrate): TSMC (TSM US / 2330 TT)’s core advanced packaging technology that densifies assembly of GPU and HBM like building blocks; currently in extreme short supply.

Edge AI (edge AI) / AI PC / AI Phone: AI applications that run on-device—directly on laptop, smartphone, or car chips—without relying on cloud servers.

Embodied AI (embodied intelligence): Putting the AI "brain" (LLMs) into a physical "body" (e.g., UBTECH Robotics (09855 HK) humanoid robots, autonomous vehicles) to enable interaction and work in the physical world.

Fine-tuning (fine-tuning): On top of pretraining, feeding industry-specific data (e.g., legal, financial, medical) so the model evolves from a "generalist" to a "domain expert."

GenAI (Generative AI, generative artificial intelligence): As opposed to traditional "recognition-type" AI, refers to AI that can autonomously create new content (text, images, video, code).

GPU (Graphics Processing Unit, graphics processor): The core chip for AI compute (e.g., NVIDIA (NVDA) H100/B200), excels at massively parallel computing.

GPU Cloud (compute cloud) / GPU-backed Debt: A novel financing model that became widespread in 2026. Startups pledge their inventory of NVIDIA (NVDA) B200/H100 chips as physical collateral to borrow from Wall Street banks.

Hallucination (hallucination): An industry term for when LLMs, due to data gaps or probabilistic errors, generate seemingly plausible but entirely incorrect, fabricated information.

Hardware & Infra: hardware and infrastructure layer.

HBM (High Bandwidth Memory, high-bandwidth memory): Ultra-high-speed memory placed adjacent to the GPU die; a key hardware element to alleviate the AI compute transfer bottleneck (the "memory" bottleneck).

Hyperscalers (ultra-large-scale cloud providers): AI's ultimate buyers—tech giants pouring capital into data centers such as Microsoft (MSFT), Amazon (AMZN) AWS, Google Cloud (GOOGL), and Meta (META).

Inference & Optimization: Efficient model training and engineering layer

Inference (推理): The second half of the AI value chain. Refers to the process after a model is trained when users use it (e.g., you asking me a question and AI generating an answer). In 2026, institutions will focus most on the decline in "inference cost."

InfiniBand (IB networking) vs RoCE: The superhighway inside AI clusters. InfiniBand is a high-end network communications protocol monopolized by NVIDIA (NVDA) (low latency, ultra-fast); RoCE is an upgraded version of traditional Ethernet (cost-effective). Their strategic battle directly impacts the earnings of network-chip names such as Broadcom (AVGO) and Marvell (MRVL).

LLM (Large Language Model, 大語言模型): Mega AI models trained on massive text corpora with tens of billions to trillions of parameters, such as GPT-4, ERNIE Bot, Zhipu GLM.

MaaS (Model as a Service, 模型即服務): The business model of foundation-model companies—rather than selling software, they expose models via cloud APIs and charge enterprises by usage.

Models & Architecture: Model and algorithm architecture layer

MoE (Mixture of Experts, 混合專家模型): The hottest architecture today (used by GPT-4, Zhipu AI (2513 HK), Minimax). Instead of one monolithic model handling everything, split it into multiple "expert sub-models" and activate the most competent ones, sharply reducing compute cost per answer.

Multimodal (多模態): AI that not only understands text but can also understand and generate images, audio, video, code, and other modalities (e.g., GPT-4o or Minimax's audio-video models).

Networking & Scaling: Compute interconnect and cluster layer

Neuromorphic Computing (類腦計算) / Spiking Neural Networks (脈衝神經網路): Next-gen chips and algorithmic architectures that simulate neuronal spiking in the human brain, aiming to deliver LLM-level or superior intelligence at ultra-low power (a few watts).

Next-Gen AI Concepts: Frontier AI concepts

NVLink / NVSwitch: NVIDIA (NVDA)'s "proprietary glue." Technology that connects tens of thousands of GPUs into a single giant chip, forming the communications substrate that fortifies NVIDIA's ecosystem moat and makes it extremely hard for AMD (AMD) to catch up.

Open-Source (開源) vs Closed-Source (閉源): A battle of approaches. The open-source camp, led by Meta (META) (Llama), gives the code away for free to the world; the closed-source camp is led by OpenAI. This shapes whether startups such as Zhipu AI (2513 HK) and Minimax build their own base models or build applications atop the open-source ecosystem.

Parameters (參數量): A metric for a model's "brain capacity." The larger the parameter count (e.g., 70B represents 70 billion parameters), the smarter the model is in theory—but compute demand grows geometrically.

Pre-training (預訓練): Feeding AI the internet's massive corpus to learn the basic rules of human language. This stage is the most cash- and compute-intensive.

Prompt Engineering (提示詞工程): Writing and refining instructions (Prompts) fed to AI to squeeze out the best possible answer.

Quantization (量化): A "model slimming" technique. It "compresses" trillion-parameter models so that, with minimal accuracy loss, they can run on ordinary Dell (DELL) servers or even on Apple (AAPL) smartphones (Edge AI).

RAG (Retrieval-Augmented Generation, 檢索增強生成): The remedy for AI's "confident nonsense (Hallucination)." Have the model retrieve accurate information from enterprise internal databases or via web search before generating the answer.

RLHF (Reinforcement Learning from Human Feedback, 人類回饋強化學習): Humans act as judges, scoring AI outputs to correct errors and align AI behavior with human values and norms.

Scaling Law (尺度定律): AI's Moore's Law—so long as data volume, compute, and model parameters increase, AI capability strengthens without bound.

Silicon Photonics (silicon photonics technology): A disruptive technology that uses "light" instead of "electricity" to transmit data between chips, sharply reducing AI server power consumption while boosting speed.

Sora-like / Text-to-Video: Text-to-video foundation models. The multimedia holy grail of the AI era and a key arena where outbound unicorns such as Minimax are intensively competing.

Synthetic Data (synthetic data): High-quality human-generated text and video are being rapidly consumed by AI; current AI is starting to use data generated by high-performance AI itself—cleaner and more logically consistent—for next-generation training.

Tokens: The basic unit for AI language processing (typically one word or Chinese character equals 1–2 Tokens). Large models are usually priced on a "per million Tokens" basis.

Training & Optimization: training and fine-tuning layer

Transformer: The underlying neural network architecture for virtually all mainstream AI (including ChatGPT); its core is the attention mechanism (Attention), enabling AI to understand context.

KGI

Investment Products and Solutions

Kevin Yin
SFC Licensed Representative (CE No. AJM689)
kevin.h.yin@kgi.com

Raymond Mok, CFA
SFC Licensed Representative (CE No. BHJ465)
raymond.mok@kgi.com

Derek Yip, CFA
SFC Licensed Representative (CE No. BSQ196)
derek.yip@kgi.com

Cusson Leung, CFA
SFC Licensed Representative (CE No. ADU276)
cusson.kt.leung@kgi.com